SPeCRA: A Stochastic Performance Comparison Routing Algorithm for LSP Setup in MPLS Networks

J. C. de Oliveira, F. Martinelli*, and C. Scoglio

Broadband and Wireless Networking Laboratory
School of Electrical and Computer Engineering
Georgia Institute of Technology, Atlanta, GA 30332
Email: {jau, caterina}@ece.gatech.edu

* Dipartimento di Informatica Sistemi e Produzione
Università degli Studi di Roma 'Tor Vergata'
00133 Roma - Italia
Email: martinelli@disp.uniroma2.it

Abstract—**In this paper, a new algorithm for dynamic routing of LSPs is proposed. While off-line algorithms are not suitable due to the necessary a priori knowledge of future LSP setup requests, our proposed algorithm, SPeCRA (Stochastic Performance Comparison Routing Algorithm), does not assume any specific stochastic traffic model and does not require any knowledge of future demands. Both features are a must for the new Internet traffic. In order to analyze SPeCRA's performance, we compare the LSP rejection ratio of SPeCRA and MIRA [1]. SPeCRA is easy to be implemented – can be implemented using only simple shortest-hop or shortest-cost algorithms, which are interesting solutions for vendors, and not as computationally heavy as other routing algorithms.**

I. Introduction

THE most promising routing algorithms proposed in the literature can be grouped into two categories: the ones based on optimal load sharing, which optimize performance indexes (e.g., bandwidth, delay) [2]; and the ones based on estimation of the current and on the forecast of the future network state [3]. These routing algorithms assume the knowledge of the stochastic characteristics of the offered traffic, which is not a realistic assumption for the current Internet.

Currently adopted routing schemes, such as OSPF and IS-IS, have the number of hops as the only metric used for routing calculations, which is not enough for QoS routing purposes. In order to introduce QoS requirements in the routing process, the widest-shortest-path (WSP) [4] and the shortest-widest path (SWP) [5] algorithms were designed. Modifications to these routing algorithms have been proposed in order to reduce complexity, such as the K shortest path, K widest-shortest-path, etc., which consider only K path options in their decisions [6]. In [7], the authors discuss two cost components of QoS routing: complexity and increased routing protocol overhead. Improvements are suggested in order to diminish this cost components, such as path pre-computation and non-pruning triggering policies. Some of these suggestions were used in our routing algorithm. Finally, another QoS routing algorithm called Minimum Interference Routing Algorithm (MIRA) is presented in [1]. MIRA tries to minimize the interference between different routes in a network for a specific set of ingress-egress nodes. The shortcomings of MIRA include its computation burden, the fact that it uses longer paths than shortest-path routing schemes, it is not able to estimate the interference effects on clusters of nodes, and its is not very likely to be implemented by vendors because of its complexity.

In this paper, we consider the problem of setting up bandwidth guaranteed LSPs in an MPLS network, where LSP setup requests arrive individually, and future requests are not a priori known. We propose SPeCRA (Stochastic Performance Comparison Routing Algorithm) to solve this problem. SPeCRA attempts to adaptively choose the best routing algorithm from a number of candidate algorithms, each of which may be suited for a different type of traffic mix. The results presented show that adaptively choosing between many different, fairly simple, algorithms results in better performance than a more complicated, computationally expensive, algorithm like MIRA.

The rest of this paper is organized as follows: the routing problem formulation and SPeCRA's description are discussed in Section II. Experimental results are analyzed in Section III. SPeCRA is applied to a commercial network and the LSP setup rejection probability is evaluated. Performance comparison between MIRA and SPeCRA are shown in Section IV. Finally, Section V includes a summary of the conclusions.

II. SPeCRA and the Routing Problem

In this section, a formulation of the LSP routing problem is presented, including explanations on how to use SPeCRA for its solution.

Assume a LSP routing scheme (e.g. shortest-path) is denoted by θ. A set of possible LSP routing schemes (e.g. shortest-path, K-shortest-path, shortest-widest-path, etc.) will be denoted by Θ. We will consider the percentage of rejected LSP setup requests in an interval to be an estimate of the probability of LSP setup rejection. This percentage depends on the LSP routing scheme adopted and also on the particular realization w of the stochastic process characterizing the traffic. In conclusion, the percentage of rejected LSP setup requests from time 0 to t can be denoted by the function $f(0,t,\theta,w)$.

If the traffic stochastic process is ergodic, the percentage of rejected LSP setup requests in the interval $[0,t]$ reaches the steady state value. This value does not depend any longer on the particular realization, but only on the routing scheme θ, and represents the LSP setup rejection probability of the network at steady state, with routing scheme θ, $p_B(\theta)$.

Suppose the stochastic process representing the traffic is stationary and let $f(0,t,\theta,w)$ be the fraction of rejected LSP setup requests in the interval $[0,t]$, with the LSP routing scheme θ and with the stochastic process realization w; then with probability 1, the limit $p_B(\theta) = \lim_{t\to\infty} f(0,t,\theta,w)$ ex-

0-7803-7632-3/02/$17.00

ists and represents the LSP setup rejection probability of the network at the steady state. The LSP setup rejection probability $p_B(\theta)$ is the objective function to be minimized. In the stationary scenario, the optimization problem is that of finding $\theta^* = \arg\min_{\theta\in\Theta} p_B(\theta)$

If the stochastic process characterizing the traffic non-stationary, which is the case of interest in this paper, the previous argument can be extended to the case of "piecewise stationary" traffic: i.e., If there exists a set $\{\tau_1, \tau_2, \cdots\}$ of time instants, which we call *switching times*, the stochastic process characterizing the traffic stationary in every time interval $[\tau_l, \tau_{l+1})$ (lth steady interval), and it becomes non-stationary at every switching time. In the following we will assume that the steady intervals are long enough that the network can be considered at steady state for a large fraction of them.

With reference to piecewise-stationary traffic, we want to design an algorithm which is able to determine, in every steady interval, the optimal LSP routing scheme in Θ. We will call $\theta^*(t)$ the optimal LSP routing scheme at time t, which leads to the minimal LSP setup rejection probability at the steady state if the system remains stationary after the time t, i.e., $\theta^*(t) = \arg\min_{\theta\in\Theta} p_B(\theta, t)$. Our objective is that of finding, at every time instant, the optimal $\theta^*(t)$.

Consider now a short time interval $[t, t+T]$, which is contained in a steady interval (namely, there exists an l such that $\tau_l \leq t < (t+T) < \tau_{l+1}$), during which no change is made in the LSP routing scheme. An estimate of $p_B(\theta, t)$ is given by the fraction of LSP setup requests rejected in such an interval: $f(t, t+T, \theta, w)$.

If we change the LSP routing scheme θ into θ', we can verify that, in general, $f(t, t+T, \theta', w) \neq f(t, t+T, \theta, w)$. Under very conservative assumptions, it is possible to prove [8] that the estimate of the order between θ and θ' in terms of LSP setup rejection probability, is more robust than the estimate of the cardinal values of the two LSP setup rejection probabilities. In fact, if there are N independent estimates of $p_B(\theta, t)$ and $p_B(\theta', t)$ taken on N different and non-overlapping intervals, the convergence rate of the estimated order to the real order is an exponential function of N and is much larger than the convergence rate of the cardinal estimates, whose variance approaches 0 with $1/N$. Such an interesting feature is used in SPeCRA, where we assume that the piecewise stationary characterizations of the traffic are denoted by SS_i,with $0 \leq i \leq I$.

Summarizing, we have a discrete and finite set $\{SS_0, SS_1, \cdots, SS_I\}$ of stationary stochastic processes, from which we can compose a non-stationary traffic by selecting any combination of SS_is (A non-stationary traffic composed by several SS_is is shown in Fig. 1). For each element SS_i, there exists a routing scheme θ_i which is optimal, within a set of possible routing schemes Θ (routing scheme θ_i leads to the minimum rejection rate $p_B(\theta_i, t)$). We do not know a priori which is the current SS_i, neither where the switching times among the SS_is are located. SPeCRA should be able to determine the optimal θ_i without knowing which SS_i is the current traffic offered to the network. The details of SPeCRA are described next.

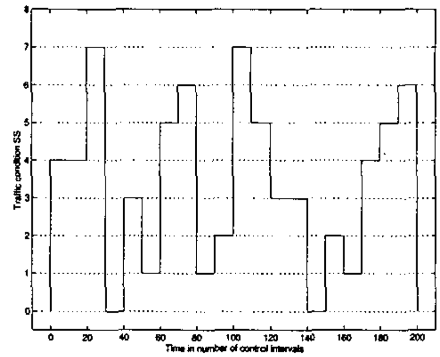

Fig. 1. Non-stationary traffic profile.

Stochastic Performance Comparison Routing Algorithm

In order to explain SPeCRA's implementation, we define an increasing sequence of time instants $t_k = kT_c$, $k = 1, 2, \cdots$, and denote the interval $[t_k, t_{k+1}]$ as the k-th control interval.

The algorithm always behaves as a homogeneous Markov chain and the optimal routing scheme is a state of the chain which is visited at the steady state with a certain probability. Aiming to reduce the chance that we could leave the state due to estimate error, we introduce a noise filter that reduces the effect of such estimate errors. Changes are less likely to happen if the adopted routing scheme (state) is "good." This is achieved by introducing a state variable Q, which reduces the changes from "good" to "bad" routing schemes, while does not reduce the changes from "bad" to "good." The algorithm is detailed below.

- *Data:*
  - The set of possible routing schemes Θ;
  - The probability function $R(\theta, \theta')$, which represents the probability of choosing θ' as candidate routing scheme when the current routing scheme is θ;
  - An initial routing scheme θ_0;
  - A time duration T_c.
- *Initialization:* Set $x_0 = \theta_0$, $Q_k = 0$ and $k = 0$.
- Iteration k:

1. Let $x_k = \theta$ be the current routing scheme and choose a set $S_k = [z_1, z_2, \cdots, z_s]$ of s candidate routing schemes, where the selection of z_i is made according to $R(\theta, z_i)$;

2. Record all the LSP setup requests arrived and ended during the interval $[t_k, t_{k+1}]$: Compute $f(t_k, t_k + T_c, z_i, w)$, $i = 1,2,\cdots,s$, the estimates of the LSP setup rejection probabilities $p_B(z_i, t)$ for each routing scheme z_i. Select $\theta' = \arg\min_{i=1,2,\cdots,s} f(t_k, t_k + T_c, z_i, w)$;

3. Choose a new routing scheme according to the estimates computed in the previous step: let $f(t_k, t_k + T_c, \theta, w)$ and $f(t_k, t_k + T_c, \theta', w)$ be the estimates of the LSP setup rejection probabilities $p_B(\theta, t)$ and $p_B(\theta', t)$ for the two schemes θ

and θ',resp ectively, in the k-th control interval.

If $f(t_k, t_k + T_c, \theta, w) - Q_k > f(t_k, t_k + T_c, \theta', w)$

$x_{k+1} = \theta'$; $Q_{k+1} = 0$;

else

$x_{k+1} = x_k$; $Q_{k+1} = [Q_k + f(t_k, t_k + T_c, \theta', w) - f(t_k, t_k + T_c, \theta, w)]/2$;

4. Set $k = k + 1$ and go to step 1.

III. Experimental Results

In this section, SPeCRA is tested through several simulations that uses an actual commercial backbone network topology, shown in Fig. 2. The topology has 11 nodes and 19 links, each link with capacity equal to 100 capacity units.

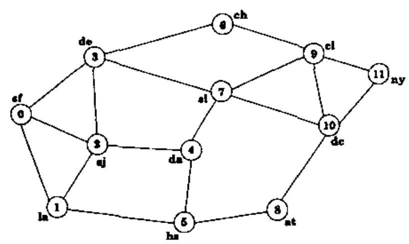

Fig. 2. US network topology for experiments.

The LSP bandwidth requests are taken to be uniformly distributed between 1 and 5 capacity units. The experiments are run for a dynamic network, in which LSPs are setup and torn down according to available bandwidth and holding times, respectively. LSP requests arrive for an ingress-egress node-pair according to a Poisson process with an average rate λ, and holding times exponentially distributed with mean $1/\mu$.

We defined 8 stationary traffic conditions: SS_0, for which $\lambda/\mu = 100$; SS_1 with $\lambda/\mu = 80$; SS_2 for which node-pair 2-11 is loaded with $\lambda/\mu = 200$, while the other nodes have much lower load: $\lambda/\mu = 50$; SS_3, for which node 6 is overloaded with $\lambda/\mu = 200$, while the other nodes have $\lambda/\mu = 70$; SS_4, for which node 8 is overloaded with $\lambda/\mu = 100$, while the other nodes have $\lambda/\mu = 60$; SS_5, for which node-pair 2-11 has $\lambda/\mu = 400$, while the other node-pairs have $\lambda/\mu = 50$; SS_6, for which $\lambda/\mu = 180$; and SS_7, with $\lambda/\mu = 150$.

For this experiment, we use the traffic profile shown in Fig. 1. The minimum interval in between two consecutive traffic conditions in the figure, during which the traffic is stationary, is equal to 200 seconds, and the control interval T_c was chosen as 20 seconds, which means a transition can happen after every 10 control intervals. The time is shown in number of control intervals, and the total simulated time is 4000 seconds.

The set of routing schemes, Θ, was chosen to be composed of three known routing algorithms (for their strength and simplicity): K-shortest-path, K-widest-shortest-path, and maximum utility path [9]. The maximum utility algorithm was implemented for K-shortest-paths. A utility value is calculated for each of the K-shortest-paths by the following function: $U(P) = \log(Abw) - \Delta h$, where Abw represents the available bandwidth along path P and Δh represents the difference between the number of hops in path P and the number of hops in the shortest-path. For each traffic matrix SS_i, the three routing algorithms were run and the probability of LSP setup rejection was calculated and compared in order to determine the best algorithm in Θ for that traffic load. This information was used to check SPeCRA's accuracy in finding the optimal routing scheme for the non-stationary traffic shown in Fig. 1. For this experiment, we chose $R(\theta, \theta')$ as equally distributed among the routing schemes. θ_0 was chosen as the shortest-path algorithm.

For each control interval, SPeCRA should find which of the three algorithms lead to the minimal rejection rate. This is done by running a short simulation for each routing scheme not currently in use for stored information about LSP arrival and tear down during the past interval. The shortest-path was the best choice when the network was generally overloaded, while the widest-shortest-path and maximum utility path algorithms were chosen in underload and local overload conditions.

Fig. 3 shows the simulation results in a illustrative manner. The sequence of the routing schemes selected by SPeCRA at each control interval T_c was translated to the sequence of SS_is for which the selected routing scheme is optimal. As we can see, the selected routing schemes coincide with the optimal solutions for the different SS_is in most of the cases run, which means that SPeCRA is indeed choosing the best routing scheme available for each traffic mix, without a priori knowledge of the nature of the mix.

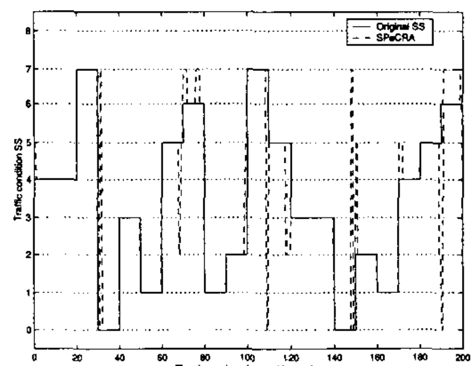

Fig. 3. SPeCRA's performance.

Even if some errors in the detection of the optimal routing scheme still remain with SPeCRA, these errors have a negligible impact on the performance, since some of the chosen routing schemes have quite similar performance in some of the eight traffic conditions.

In Fig. 4, we observe the LSP setup rejection rate for every traffic condition SS_i when using shortest-path only and when using SPeCRA to select the most appropriated routing scheme. We observe that SPeCRA selects a much better routing scheme when SS_0, SS_1, SS_3, and SS_4 are the traffic conditions in use.

For the non-stationary traffic in Fig. 1, SPeCRA leads to over 25% performance improvement when compared to using the shortest-path algorithm as the fixed choice for routing scheme. A non-stationary traffic that has more components such as SS_0, SS_1, SS_3, or SS_4 (e.g., suppose we randomly choose 20 traffic conditions from this set to compose a non-stationary traffic), SPeCRA would lead to even higher improvement regarding LSP setup rejection rate. It is important to notice that SPeCRA does switch between routing schemes in every control interval (which would be costly). The algorithm will only switch when the new routing scheme results in a better performance.

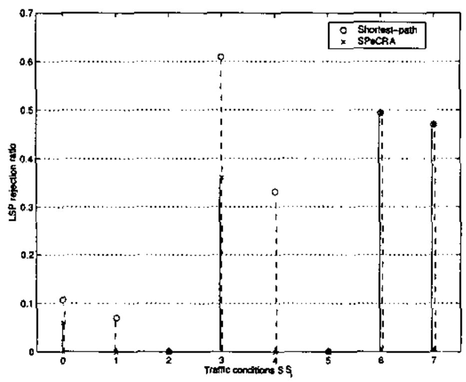

Fig. 4. LSP rejection ratio for shortest-path and SPeCRA.

IV. Performance Evaluation

In this section, we will evaluate SPeCRA's performance by comparing its LSP setup rejection ratio with MIRA's(M inimum Interference Routing Algorithm) [1]. MIRA is a dynamic algorithm that routes bandwidth guaranteed tunnels (LSPs). The algorithm is based on the idea that a new LSP must be routed in a path that does not "interfere too much" with routes that could be taken by future demands. Such path must be found without any a priori knowledge of the future demands. A set of ingress-egress pairs is selected and the algorithm is run in order to select the "critical links" [1]. For every LSP setup request arrived, maxflow values are calculated, critical links are discovered (heavy computations), weights are computed in order to be used by the weighted shortest-path algorithm. All links which have residual bandwidth below a certain threshold are eliminated, resulting in a smaller network. The weighted shortest path is then calculated, the route is finally chosen, the LSP is routed, and the residual capacities are updated. Even though critical links do not need to be discovered on every LSP arrival, unless the threshold is chosen to be 1, the algorithm is still heavy computationally. Moreover, in an "all ingress-egress network" (every node on the network can be a ingress or egress node, such as the US topology in Fig. 2), the algorithm would not be able to avoid interference as much as on the examples presented on the paper, which would make its behavior similar to the shortest-path algorithm. SPeCRA may still be able to perform better than shortest-path in this scenario.

In order to compare the performance of each method, we did not implement MIRA (as the implementation's accuracy could be questionable) but instead run SPeCRA for the same network, with the same traffic conditions described in [1]. The performance of each algorithm is measured by the number of LSPs rejected by the algorithm.

Consider the network shown in Fig. 5, which is the same network considered in [1]. The network is composed by 15 nodes and 28 links, and the only ingress-egress pairs were chosen as shown. The capacity of the thin-line links is 12 units and 48 units for the thick-line links. All the links are bidirectional (acting like two unidirectional links of that capacity). The LSP setup requests arrive randomly, at the same average rate, for the ingress-egress pairs. The bandwidth request are chosen to be uniformly distributed between 1 and 3 capacity units.

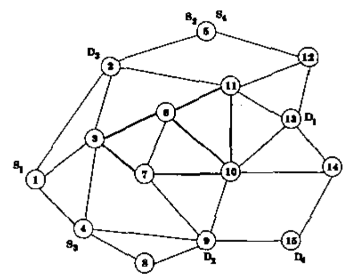

Fig. 5. Network topology for performance simulations.

The Static Case

In the first set of experiments, all the LSPs are assumed to be static (once setup they stay in the network forever and are not torn down). The link capacities are now scaled by 100 in order to fit a larger demand. 5000 LSP setup requests are generated and the number of rejections in 20 individual executions of the program is observed. The results are shown in Fig. 6.

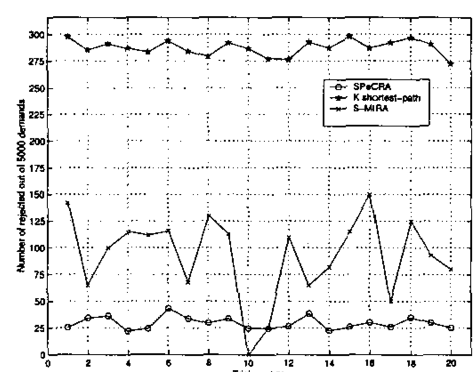

Fig. 6. Static case: number of LSP setup requests rejected in 20 trials.

We observe that the K-shortest path algorithm has a larger rejection rate due to the fact that in most of the cases there is only a few options for the shortest path between the ingress-egress pairs (S_1-D_1 has one option, S_2-D_2 has three options, S_3-D_3 has 2 options, and S_4-D_4 has only one option). When more than one option is available, the algorithm chooses the one with more residual bandwidth, and in that case, tries to distribute the load over the paths with the same number of hops. Another algorithm considered by SCA in this simulation is the K-widest-shortest path. In order to simplify computational load, the algorithm first selects the 10 shortest paths and then selects the one among them with the larger residual bandwidth (the number 10 was chosen balancing complexity and effectiveness). Similarly, the maximum utility algorithm chooses, among the 10 shortest paths, the one with the larger utility coefficient. SCA chooses the algorithm with smaller rejection rate, and the results are very satisfactory. A low rejection ratio is achieved with a "not so high" computational cost. SCA is adaptive to the traffic profile and also does not need any information about future demands.

It is important to stress that we did not reproduce L-MIRA results in Fig. 6, where only S-MIRA values were plotted. With the objective of making MIRA less computationally intensive, L-MIRA was developed with an approximation for an NP-hard problem (LEX-MAX) [1]. The approximation taken is to "assume that the ordering of the maxflow values, after the current demand is routed, will be the same as the current ordering of the maxflow values." This is accurate considering that the demand to be routed is small, as the authors stress in the paper. We believe that this approximation does not hold for the traffic offered in this static case, which is not small. Therefore, we only consider S-MIRA implementation for the performance experiments.

The Dynamic Case

With this experiment we will evaluate the performance of MIRA and SPeCRA in a dynamic scenario. LSPs arrive between every ingress-egress pair according to a Poisson process with average rate λ and holding times exponentially distributed with mean $1/\mu$. In order to compare our results to MIRA's, we will use a fixed value for $\lambda/\mu = 150$. Bandwidth demands are again uniformly distributed between 1 and 3 capacity units. We now scale the link capacities in Fig. 5 by 10 instead of 100. The rejection ratio is calculated over 1000000 LSP setup requests. Fig. 7 shows the results. We observe that SPeCRA obtains a better performance than MIRA.

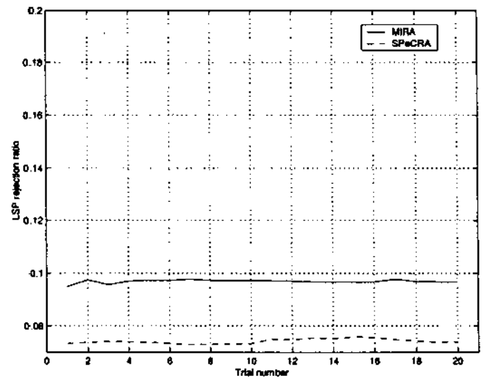

Fig. 7. Dynamic case: LSP rejection ratio in 20 trials.

V. CONCLUSIONS

In this paper we have proposed a new routing algorithm, SPeCRA, which does not assume any stochastic model to represent the incoming traffic. Given a set of routing schemes and some periodic network state information, SPeCRA is able to select the best routing scheme in the set for current traffic. In order to analyze SPeCRA's performance, we run simulations of the same network topology and traffic described in [1] and compare the LSP rejection ratio obtained by SPeCRA and MIRA. We observe that SPeCRA outperforms MIRA in the static and also dynamic cases. The set of routing schemes chosen for SPeCRA's simulations is comprised of simple shortest-hop or shortest-cost algorithms, which are interesting solutions for vendors that would rather not implement a complicated algorithm. SPeCRA is easy to be implemented and not as computationally heavy as other routing algorithms. As future work, the authors are investigating how to speed up the online simulations and also the use of other routing schemes.

ACKNOWLEDGMENTS

This work was supported by CAPES (The Brazilian Ministry of Education Agency) and NASA Goddard and Swales Aerospace under contract number S11201 (NAS5-01090).

REFERENCES

[1] K. Kar,M. Kodialam, and T. V. Lakshman, "Minimum Interference Routing of Bandwidth Guaranteed Tunnels with MPLS Traffic Engineering Application," *IEEE Journal on Selected Areas in Communications*, vol. 18, no. 12, pp. 2566–2579, December 2000.

[2] J. W. Roberts, "Traffic Theory and the Internet ," *IEEE Communications Magazine*, vol. 39, no. 1, pp. 94–99, January 2001.

[3] U. Ahlfors, U. Korner, and M. Pioro, "Towards Efficient Call Admission Control for State-Dependent Routing in Multirate Networks," in *Proceedings of 1999 IEEE International Conference on Control Applications*, 1999, vol. 2, pp. 1607–1613.

[4] R. Guerin, A. Orda, and D. Williams, "QoS Routing Mechanisms and OSPF Extensions," in *Proceedings of 2nd Global Internet Miniconference (joint with Globecom'97)*, Phoenix, AZ, November 1997.

[5] Z. Wang and J. Crowcroft, "Quality-of-Service Routing for Supporting Multimedia Applications," *IEEE Journal on Selected Areas in Communications*, vol. 14, no. 7, pp. 1288–1234, September 1996.

[6] D. Eppstein, "Finding the k shortest paths," *SIAM J. Computing*, vol. 28, no. 2, pp. 652–673, 1998.

[7] G. Apostolopoulos, R. Guerin, S. Kamat, and S. K. Tripathi, "Quality of Service Based Routing: A Performance Perspective," in *Proceedings of SIGCOMM'98)*, 1998.

[8] L. Dai and C. H. Chen, "Rates Convergence of Ordinal Comparison for Dependent Discrete Event Dynamic Systems," *Journal of Optimization Theory and Applications*, vol. 94, no. 1, pp. 29–54, July 1997.

[9] S. Oueslati-Boulahia and E. Oubagha, "A Comparative Study of Routing Algorithms for Elastic Traffic in a Multiservice Network," in *Proceedings of ITC Specialist Seminar on IP Traffic*,Mo nterey, CA, USA, 2000.