

Evaluation of Resource Reservation Policies for Deterministic Services in Multi-Service Packet Networks

Khaled M. F. Elsayed, Amr S. Ayad, and Mahmoud T. El-Hadidi
khaled@ieee.org, asaad@salec.com, hadidi@mailier.eun.eg

Department of Electronics and Communications Engineering, Cairo University, Giza, Egypt
12613

We address the issue of reserving resources at packet switches along the path of connections requiring a deterministic end-to-end delay bound. The switches are assumed to schedule outgoing packets using either the Packet-by-Packet Generalized Processor Sharing (PGPS) or the Rate-Controlled Earliest-Deadline First (RC-EDF) scheduling disciplines. We propose a set of resource reservation policies that are used to map the end-to-end delay requirement into a local rate or local deadline to be reserved at each switch. We present simulation results to evaluate the performance of these resource allocation policies. For PGPS it is found that in general network topologies when shortest-path first routing is used, that uniform allocation performs better than other policies. In the case of RC-EDF, it is shown that non-uniform delay allocation policies provide resilient and superior performance over uniform allocation of resources. This is especially clear in networks with improper capacity planning or heterogeneous links.

1. INTRODUCTION

One of the main promises of multi-service packet networks is to provide applications with Quality-of-Service (QoS) guarantees, such as maximum packet delay and packet loss ratio. Handling the variety in QoS requirements of different applications requires the network to use a mechanism for serving packets from different applications according to their granted QoS level. Many scheduling disciplines have been proposed in the literature to implement such mechanism (see [2, 3, 6, 7, 11]). Each scheduling discipline requires algorithms for performing connection admission control (CAC) and resource reservation. Most of the work in the literature assumes uniform resource reservation. In this paper, we propose non-uniform resource allocation policies for the case of the PGPS and RC-EDF scheduling disciplines and connections requiring a hard (deterministic) bound on end-to-end delay. The paper addresses the following problems:

- How to map the end-to-end delay requirement of a connection into a local resource requirement to be reserved at each switch along the connection's path?
- How to divide the resource requirement among the switches on the connection's path? We explore whether using policies that take the capacities of the switches and/or their loading into account (instead of using uniform allocation) would provide performance gains.

There have been several trials to handle the problem of resource reservation and connection admission control for packet networks providing deterministic service. In [10], it is found, that with the packet loss probability as the QoS metric, there is little difference in the performance of allocation policies in the regime of applications with low loss requirements. However, for applications which tolerate large packet loss or alternate QoS metrics, QoS allocation policies differ significantly in their performance.

Ref. [4] presents a general framework for admission control and resource reservation for multicast sessions. The problem of admission control is decomposed into several subproblems that include: the division of end-to-end QoS requirements into local QoS requirements, the mapping of local QoS requirements into resource requirements, and the reclaiming of the resources allocated in excess.

Recently, Ref. [9] considers optimal partitioning of QoS requirements for unicast paths and multicast trees. The problem is mapped to a constrained optimization problem and heuristics are used to find an approximate solution for the optimization problem. While the approach is very general, the problem with it is the assumption of the existence of a cost function. The cost function is assumed to be known and is static in nature. This is not a realistic assumption since the cost function should depend on the current network loading state and the QoS objectives of the connections passing through a particular switch.

In Ref. [1], a study of resource allocation for PGPS based scheduling in packetized voice networks is presented. It is shown that non-uniform allocation provides a performance gain. However, it only handles VoIP traffic in a particular network topology.

The rest of this paper is organized as follows: Section 2 presents the system model and details of the PGPS and RC-EDF scheduling disciplines. Section 3 includes description of the CAC algorithm and details about the proposed resource reservation policies that are applicable when using PGPS and RC-EDF scheduling. Section 4 presents simulation results, and section 5 concludes the paper.

2. SYSTEM MODEL AND DETAILS OF THE PGPS AND RC-EDF SCHEDULING DISCIPLINES

We consider a network with arbitrary topology represented by a graph $G=(V, E)$ where V is the set of network nodes representing packet switches containing the schedulers and E is the set of edges connecting these nodes. An edge $e \in E$ has an associated capacity C^e expressed in bits/sec. The network handles connections with arbitrary QoS requirements; however, we are only concerned with the class of delay-guaranteed connections requesting a deterministic end-to-end delay bound. Such connections are usually assigned a fixed portion of bandwidth. To optimize the resource usage, this reserved bandwidth may be shared with best-effort traffic when no packets from the class of delay-guaranteed connections exist. Also, packets from delay-guaranteed connections can preempt best-effort packets. Under these assumptions, we can isolate the delay-guaranteed connections from other connections and analyse them separately. Each switching node has a packet scheduler for scheduling packets from the delay-guaranteed connections. The scheduling discipline is either PGPS or RC-EDF.

Each connection f from the delay-guaranteed class is shaped by a token bucket traffic shaper and its traffic is specified by either the (σ_f, ρ_f) descriptor or the (σ_f, ρ_f, c_f) descriptor, where σ_f is the maximum burst size (in the appropriate units of bits or octets), ρ_f is the

sustained rate (in bits/sec or octets/sec), and c_f is the peak rate (in bits/sec or octets/sec). For the (σ, ρ, c) model, we further define the parameter $a = \sigma / c$, which is the time period the traffic can burst with rate c . For the (σ, ρ) model, we have $c \rightarrow \infty$, $a \rightarrow 0$, but $ac \rightarrow \sigma$.

2.1. The PGPS Scheduling Discipline

In packet-by-packet generalized processor sharing (PGPS) scheduling [2, 11] and its variants, each packet is labeled with a service tag and packets are then served in the ascending order of service tags assigned to them. Connections can be associated with service weights, and they receive service in proportion to this weight whenever they have data in the queue. The PGPS discipline can be used to provide bandwidth and delay bounds to connections with QoS guarantees as well as fair resource sharing for best-effort connections.

Noting that PGPS allocates a minimum service rate to each connection, Parekh and Gallager [11] proved an important bound on the worst-case end-to-end delay experienced by a connection traversing a series of PGPS schedulers. It was shown that if a connection f traverses a path of K_f PGPS switches and has traffic characteristics conforming to a leaky bucket with a maximum burst size of σ_f bits and a long term average rate of ρ_f bps, then an upper bound on the end-to-end delay is guaranteed for each packet from connection f by reserving a certain service rate at each switch along the connection's path. The upper bound on delay is given by:

$$D_f^i \leq \frac{\sigma_f - l^i}{g_f} + \sum_{j=1}^{K_f-1} \frac{\max_{n \in \{1, \dots, i\}} l^n}{g_f^j} + \frac{l^i}{g_f^{K_f}} + \sum_{j=1}^{K_f} \alpha^j \quad (1)$$

where D_f^i is the upper bound on the end-to-end delay encountered by the i^{th} packet from connection f , where the packets are numbered starting with the first packet to arrive in a busy period, l^i is the length of the i^{th} packet in bits, g_f^j is service rate of connection f packets at switch j , $g_f = \min_j g_f^j$, and α^j is the sum of propagation delay and maximum transmission delay over link j (the maximum transmission delay is obtained by dividing maximum packet length at switch j by link capacity). Assuming packet size is of fixed length L (for example in ATM networks), equation (1) is modified to:

$$D_f \leq \frac{\sigma_f - L}{g_f} + \sum_{j=1}^{K_f} \frac{L}{g_f^j} + \sum_{j=1}^{K_f} \alpha^j \quad (2)$$

Equation (2) shows that larger values of burst size leads to a stronger dependency of the upper delay bound on the minimum allocated rate g_f . Equations (1) and (2) are only valid when the following conditions are met at each switch j , $j \in \{1, \dots, K_f\}$.

1-**The stability condition:** This condition requires that:

$$\sum_{k=1}^{N_j} \rho_k^j \leq C^j \quad (3)$$

where N_j is the number of accepted connections at switch j and ρ_k^j is the average rate of connection k at switch j . The stability condition is necessary for all scheduling disciplines and is not specific to PGPS scheduling.

2-The PGPS schedulability condition: The schedulability condition guarantees that a scheduler will not make violations and will, therefore, honor the QoS commitment made during the connection set-up phase. For PGPS the schedulability condition requires that:

$$\sum_{k=1}^{N_j} g_k^j \leq C^j \quad (4)$$

2.2. The RC-EDF Scheduling Discipline

The operation of an EDF scheduler is described as follows: a deadline is assigned to each newly arriving packet from connection f . The deadline is computed as the sum of the arrival time of the packet and the local delay bound reserved for connection f at this scheduler. The scheduler serves packets in the ascending order of their deadlines.

Maintaining the delay guarantee made to connection f is equivalent to having all the packets belonging to it transmitted completely before their assigned deadlines. Consequently, all packets from connection f do not get delayed beyond the delay bound reserved for connection f at this scheduler. We denote the case in which a packet misses its deadline, i.e. not transmitted completely before its deadline, as a case of *violation*.

The conditions under which a single EDF scheduler operates without violations are:

- 1- **The stability condition:** This is the same as defined in section 2.1.
- 2- **The RC-EDF schedulability condition:** The necessary and sufficient schedulability condition for EDF scheduling and token-bucket shaped traffic is provided by [5]. The results are obtained by applying Theorem 1 in [8] to deduce the schedulability conditions for token-bucket traffic models. The analysis in [5] assumes the use of a preemptive EDF scheduler which is equivalent to the use of negligible packet transmission time (which is typical in the case of networks with small packet size and high speed links, e.g. ATM networks). For negligible packet transmission time, Ref. [5] defines the function $F(t)$ as:

$$F(t) = C t - \sum_{j \in N} A_j^*(t - d_j) \quad (5)$$

where C is the data rate of the link serving the EDF scheduler in bits per second, N is the set of connections passing through the scheduler, d_j is the delay bound reserved for connection j in seconds, and $A_j^*(t)$, in bits, is the traffic-constraint function on the traffic arrivals from connection j up to time t . In [5], it is shown that the *necessary and sufficient* schedulability condition of an EDF scheduler is equivalent to verifying that:

$$F(t) \geq 0 \quad \forall t \geq 0 \quad (6)$$

We will be dealing with the case in which $A_j^*(t)$ represents token-bucket traffic models. For

the (σ, ρ, c) model the function $A^*(t)$ is given by $A^*(t) = \begin{cases} ct, & 0 \leq t \leq a \\ \sigma + \rho(t - a), & t > a \end{cases}$, whereas for the

(σ, ρ) model it is given by $A^*(t) = \sigma + \rho t, \forall t \geq 0$, and for both models $A^*(t) = 0, \forall t < 0$.

We assume that the scheduler at each switch has full information about the traffic descriptor of all passing flows. Because of the fact that after flow f passes through the first switch, the traffic is no longer specified by (σ_f, ρ_f) or (σ_f, ρ_f, c_f) after the first scheduler. We assume the schedulers at each switch are rate-controlling schedulers, i.e. they reshape the traffic of a certain flow f to its original form. Since the schedulers use the EDF scheduling discipline and traffic re-shaping is performed at each switch, we have a *Rate-Controlled EDF* scheduling discipline.

3. CAC ALGORITHMS AND RESOURCE RESERVATION POLICIES

3.1. The CAC Algorithm

The CAC operation assumes that the connection is established using a setup protocol such as RSVP or ATM Q.2931 signalling. The operation of this protocol proceeds as follows: the calling party wishing to establish a connection f , sends a SETUP message to the called party, including the connection's traffic characteristics (σ_f, ρ_f, c_f) , and the required end-to-end delay bound D_f . This message travels over K_f switches belonging to the path P_f selected for the connection by the routing algorithm in use. First, each switch along the path checks the stability condition. The second test compares the value of the end-to-end delay D_f requested by the incoming connection with the value of the total transmission and propagation delay along the connection's path. If the value of the required end-to-end delay is smaller, the connection is rejected. The next test is to verify that the value of the required end-to-end delay of the incoming connection is not less than the minimum end-to-end delay bound that the network can guarantee to the incoming connection. On passing the previous tests successfully, an allocation policy is used to map the required end-to-end delay into a local resource requirement to be reserved at each switch.

3.2. PGPS Resource Reservation Policies

In this section, we derive the reserved rate for a connection according to the various proposed policies.

Even Policy (EVEN/PGPS): The reserved rates are the same at all switches, i.e.

$$g_f^i = g_f \quad \forall i \in \{1, \dots, K_f\} \quad (7)$$

Substituting in (2), after converting it to an equality to reserve the least amount of resources required for meeting the delay bound of connection f , we get:

$$g_f^i = g_f = (\sigma_f + (K_f - 1)L) / \left(D_f - \sum_{j=1}^{K_f} \alpha^j \right) \quad \forall i \in \{1, \dots, K_f\} \quad (8)$$

Capacity Proportional Policy (CP/PGPS): The reserved rate at a certain switch is proportional to the switch capacity, i.e.

$$g_f^i = \eta_f C^i \quad \forall i \in \{1, \dots, K_f\} \quad (9)$$

where $\eta_f = \text{Constant}$ for the path and connection f parameters. Substituting in (2), after converting it to an equality to reserve the least amount of resources required to meet the delay bound and solving for η_f , we get:

$$g_f^i = \left(\frac{\sigma_f - L}{C} + \sum_{j=1}^{K_f} \frac{L}{C_j} \right) / \left(D_f - \sum_{j=1}^{K_f} \alpha^j \right) C^i \quad \forall i \in \{1, \dots, K_f\} \quad (10)$$

where $C = \min_j C^j$.

Remaining Capacity Proportional Policy (RCP/PGPS): Define $R^i = C^i - \sum_{k=1}^{N_i} g_k^i$ to be the remaining capacity at switch i , where N_i is the number of connections currently allocated to switch i , and g_k^i is the rate allocated to connection k at switch i . The reserved rate to connection f at a certain switch is proportional to the remaining capacity of the switch:

$$g_f^i = \eta_f R_f^i \quad \forall i \in \{1, \dots, K_f\} \quad (11)$$

where η_f is constant for the path and connection f parameters as long as no other connections are accepted at any of the switches along the connection's path during connection setup phase. Define R_f as $\min R_i, i \in \{1, \dots, K_f\}$ and substituting in equation (1) (after converting it to an equality to reserve the least amount of resources required for meeting the delay bound) and solving for η_f , we get:

$$g_f^i = \left(\frac{\sigma_f - L}{R_f} + \sum_{j=1}^{K_f} \frac{L}{R_f^j} \right) / \left(D_f - \sum_{j=1}^{K_f} \alpha^j \right) R_f^i, \quad \forall i \in \{1, \dots, K_f\} \quad (12)$$

3.3. RC-EDF Resource Reservation Policies

The non-even policies, proposed for the RC-EDF schedulers, are based on the assumption that each scheduler initially reserves the tightest possible delay value for the incoming connection. Subsequent relaxation of this reservation by redistributing the excess end-to-end delay on the schedulers according to some criterion is then performed. Let D_f^* be the minimum achievable end-to-end delay bound obtained by having each scheduler reserve the tightest possible delay value. The excess end-to-end delay is defined as:

$$\bar{D}_f = D_f - D_f^* \quad (13)$$

Even Policy (EVEN/EDF): We use the even policy as a reference policy against which other policies may be compared. In the EVEN/EDF policy, all schedulers are required to reserve the same amount of delay, hence:

$$d_f^i = \frac{D_f}{K_f} \quad (14)$$

Even distribution of excess delay (DYNEVEN/EDF): This policy is the one suggested in [3]. The delay bound formula is given by:

$$d_f^i = d_f^{i^*} + \frac{\bar{D}_f}{K_f}, \forall i \in \{1, \dots, K_f\} \quad (15)$$

Capacity proportional distribution of excess delay (DYNCP/EDF): In this policy, the excess delay is distributed in inverse proportion to the switch link capacity. Therefore, switch i , will be assigned a portion of the excess delay in proportion to inverse of its link capacity to the sum of inverse capacities of all switches in the path. The delay bound formula is given by:

$$d_f^i = d_f^{i^*} + \frac{\bar{D}_f / C^i}{\sum_{j=1}^{K_f} 1/C^j}, \forall i \in \{1, \dots, K_f\} \quad (16)$$

Remaining-delay proportional distribution of excess delay (DYNRDP/EDF): In this policy, the excess delay is proportional to the minimum delay bound that the scheduler can guarantee to the incoming connection. The delay bound formula is given by:

$$d_f^i = d_f^{i^*} + \frac{\bar{D}_f}{D_f^*} d_f^{i^*}, \forall i \in \{1, \dots, K_f\} \quad (17)$$

$$\Rightarrow d_f^i = d_f^{i^*} \frac{D_f}{D_f^*}, \forall i \in \{1, \dots, K_f\} \quad (18)$$

4. PERFORMANCE ANALYSIS

A suitable performance metric for evaluating the performance of the policies is the connection blocking probability of the network. It is difficult to analytically obtain the connection blocking probability for general network topologies with different types of traffic being offered to the network. Therefore, simulation is used in order to compare the different policies for this case.

In our simulation, connections are generated according to a Poisson process with an average arrival rate of λ , and their duration is exponentially distributed with a mean $1/\mu$. The value $\rho = \lambda / \mu$ characterizes the offered traffic load (in Erlangs) at the network under consideration. We fix μ to be 1 unit. For each simulation, we generate a total of between 200,000 and 1 million connections (depending on traffic load and network topology). Each experiment is repeated 5 times with different seeds to obtain confidence intervals (we do not show the confidence intervals in the results as they were generally very narrow except for blocking probability values less than $1e-5$). The blocking probability is computed as the number of blocked connections divided by the number of generated connections.

We use the (σ, ρ) traffic model where realistic values of traffic descriptors are used and are selected as follows [5]:

- Average rate $\rho = 10^m$ Kbps, m is uniformly distributed on $[0, 3]$. This gives a mean average rate of 144.62 Kbps
- Burst size $\sigma = y * r$ kbits, y is uniformly distributed on $[0.5, 1.3]$. This gives a mean burst size of 130.16 Kbits
- Delay bound $= 50 * 10^s$ msec, s is uniformly distributed on $[0, 1.52]$

This range of generated traffic patterns include a typical MPEG video source with average rate 518.4 kbps, and burst size 576 kbits. It also includes a typical packetized voice source with average rate 10 kbps, and burst size 8 kbits.

We obtain results for two general networks used extensively in network performance studies. The first network topology is the COST239 network and the second one is the NSFNET topology (both of which are shown in *Figure 1*). Shortest-path-first routing is used where the link cost for all links in the network is equal to one (in other words the path selected is the path with the minimum number of hops between the source and destination). In case there are more than one shortest path, only one of these paths are selected (i.e. there is no load balancing among equal or non-equal cost multiple paths). The traffic is symmetric and divided equally among all possible source destination pairs. We evaluate the performance for two link configurations: in the first configuration all links have the same capacity of 34 Mbps (uniform links), whereas in the second configuration link capacities are randomly chosen such that overall capacity in the network is equal to the uniform links case.

Given the above configuration, we run the simulation for different loading factors for both the uniform capacity and random capacity configuration using the COST239 topology. The results are shown in *Figure 2* and *Figure 3* for the uniform capacity and random capacity settings respectively. The traffic load is varied from 2 to 2048 Erlangs. We obtain the performance of all policies in part (a) of the graph, and compare the best two performing policies of the PGPS and RC-EDF families in part (b) of the graph.

It is clear from the figures, that the set of policies belonging to PGPS or EDF family of scheduling disciplines have a similar trend/behavior. In general the PGPS based policies perform significantly better than RC-EDF policies in light load, whereas RC-EDF based policies start to be better at higher load. This is in conformance with the known fact that the delay bound obtained by PGPS is known to be conservative. Therefore, at light load, PGPS performance is not affected since there is ample capacity to use. However, when the network load increases, the performance of PGPS starts to deteriorate because of the previously mentioned nature of its delay bound.

Also, we can see that for PGPS, the EVEN policy is the best performing policy. For RC-EDF the DYNCP or DYNEVEN are the best performing policies. Also, in the random capacity case, the dynamic EDF policies perform much better than the EVEN policy, which illustrates the need for dynamic policies when EDF is used in heterogeneous networks.

Similar to the setup used for the COST239 network, we perform simulation on the NSFNET topology for uniform and random capacity setting. Similar behavior is observed as in the COST239 network with the exception that the difference between the PGPS policies and EDF policies in light load is not as big as in the COST239 network. (We do not show the graphs for the NSFNET topology due to space restrictions.)

We can conclude that PGPS combined with the EVEN policy for rate allocation is the best performing policy except under heavy load. Moreover, in RC-EDF there is extra complexity needed to reshape the traffic in each switch. Therefore, it can be suggested that PGPS with

EVEN or CP (which are also the easiest policies to implement) outperform all other policies in networks with complex topologies and shortest-path-first routing.

5. CONCLUSIONS

This paper discussed the use of non-even resource reservation policies when performing resource reservation in order to provide bounded delay service in a multi-service network. We have proposed a set of dynamic resource reservation policies for a network with general topology and dynamic heterogeneous traffic when PGPS or RC-EDF scheduling is used.

We have seen through simulation results on various topologies that the following holds true:

- For RC-EDF policies, the dynamic policies such as DYNEVEN and DYNCP perform much better than EVEN policy.
- The EVEN policy for PGPS is a very resilient policy specially for networks that are well-designed. Similar performance is obtained by the CP/PGPS policy, which can be useful in cases when proper network planning/routing is not performed. Luckily, among all policies these two policies are the easiest to implement.
- In most cases at light load, PGPS policies perform better than RC-EDF. However, for some combinations of network topology and loading state (usually at high load), RC-EDF policies perform better than PGPS policies. So, picking which policy to use is dependent upon the network topology and traffic characteristics. Given the complexity of RC-EDF (reshaping and the calculation needed by the admission process), we highly recommend using PGPS (EVEN or CP) policies for multi-services networks providing end-to-end guarantees on the delay.

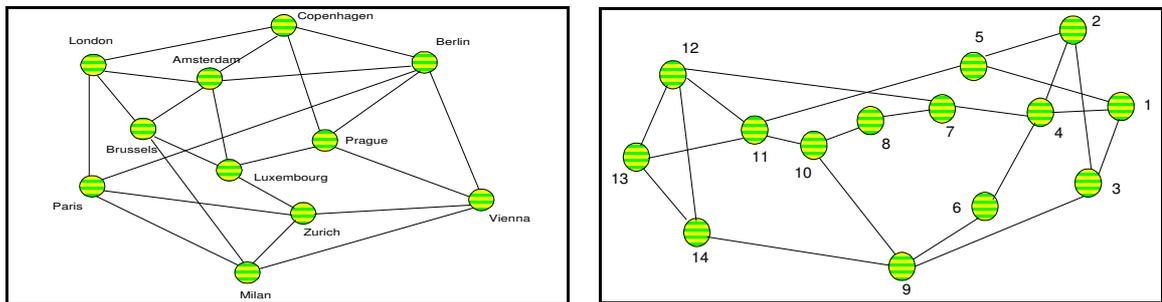
It should be noted however that using non-even policies would, in general, involve the exchange of periodic link state packets containing more information about the loading state of network links. This may cause more processing burden on the control processors of the network switches specially when employing more complex routing algorithms.

This paper has only considered the case of deterministic end-to-end delay bounds as a QoS metric, our work can be extended by considering statistical end-to-end delay bound or other QoS metrics such as bandwidth or loss bounds. Another important area of research is integration with QoS routing for selecting paths that maximizes acceptance ratio of the traffic.

REFERENCES

- [1] M. Baldi and F. Risso, Efficiency of Packet Voice with Deterministic Delay, *IEEE Comm. Magazine*, pp. 170-177, May 2000.
- [2] A. Demers, S. Keshav, and S. Shenker, Design and Analysis of a Fair Queuing Algorithm, *Proceedings of ACM SIGCOMM'89*, Austin, 1989.
- [3] D. Ferrari and D. Verma, A Scheme for Real-Time Channel Establishment in Wide-Area Networks. *IEEE JSAC*, vol. 8, no. 4, pp. 368-379, April 1990.
- [4] V. Firoiu and D. Towsley, Call Admission and Resource Reservation for Multicast Sessions, *Proceedings of IEEE INFOCOM'96*, pp. 94 – 101, March 1996.
- [5] V. Firoiu, J. Kurose, and D. Towsley, Efficient Admission Control of Piecewise Linear Traffic Envelopes at EDF schedulers, *IEEE/ACM Transactions on Networking*, Vol. 6, No. 5, pp. 558-570, Oct. 1998.
- [6] L. Georgiadis, R. Guerin, V. Peris, and K. Sivarajan, Efficient Network QoS Provisioning Based on Per Node Traffic Shaping, *IEEE/ACM Transactions on Networking*, Vol. 4, pp. 482 -501, Aug. 1996.

- [7] S. J. Golestani, A Self-Clocked Fair Queuing Scheme for Broadband Applications, Proceedings of IEEE INFOCOM' 94, pp. 636-646, Jun. 12-16, 1994.
- [8] J. Liebeherr, D. E. Werge, and D. Ferrari, Exact Admission Control for Networks with a Bounded Delay Service, IEEE/ACM Transactions on Networking, Vol. 4 No. 6, pp. 885-901, 1996.
- [9] D. H. Lorenz and A. Orda, Optimal Partitioning of QoS Requirements on Unicast Paths and Multicast Trees, IEEE/ACM Transactions on Networking, Vol.. 10 No. 1, pp. 102-114, 2002.
- [10] R. Nagarajan, J. Kurose, and D. Towsley, Local Allocation of End-to-End Quality of Service in High-Speed Networks. Proceedings of IFIP Workshop on the Performance Analysis of ATM Systems, Martinique, Jan. 1993.
- [11] K. Parekh and R. G. Gallager, A Generalized Processor Sharing Approach to Flow Control in Integrated Services Networks: The Multiple Node Case, IEEE/ACM Transactions on Networking, vol. 2, pp. 137-150, 1994.



(a) The COST239 Network Topology (b) The NSFNET Network Topology
 Figure 1. Network topologies of the COST239 network and NSFNET

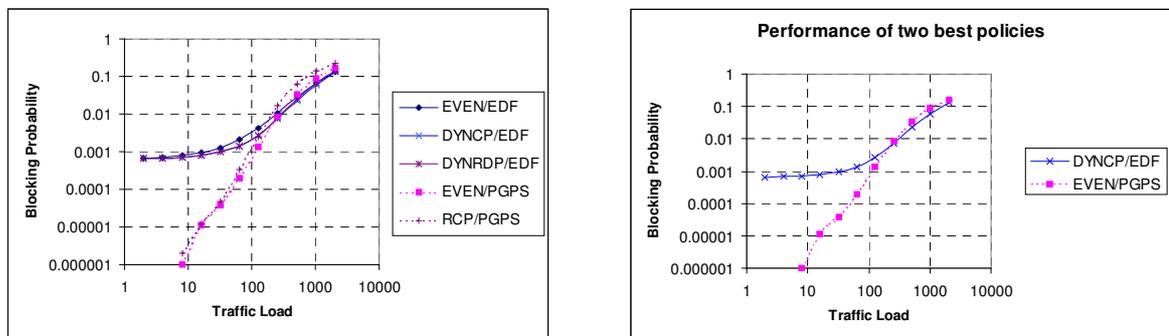


Figure 2. Effect of traffic load on the blocking probability when using the proposed policies in the COST239 network with uniform capacity and SPF routing

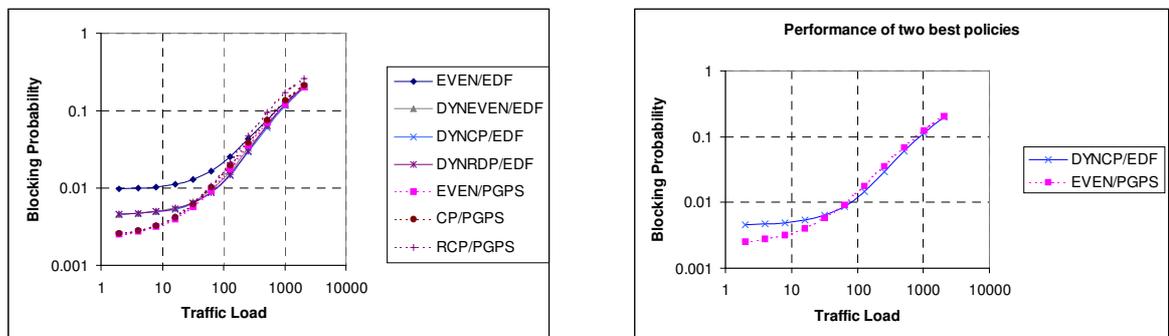


Figure 3. Effect of traffic load on the blocking probability when using the proposed policies in the COST239 network with random capacity and SPF routing