A Framework for End-to-End Deterministic-Delay Service Provisioning in Multiservice Packet Networks

Khaled M. F. Elsayed, *Senior Member, IEEE*

***Abstract*— The problem of providing end-to-end delay guarantees for deterministic-delay services in multi-service packet networks is addressed through a combination of dynamic resource reservation and routing. Our model is based on using rate-controlled earliest-deadline-first (RC-EDF) for providing hard bounds on end-to-end delays. With RC-EDF a certain delay bound has to be allocated for a connection at each node in the selected path. The most commonly used resource reservation policy is uniform allocation which is based on dividing the end-to-end delay bound equally among the nodes in the selected path. This simple allocation policy could lead to non-uniform resource loading and subsequently lead to high blocking rates. Moreover, the most commonly used routing method is shortest-path first routing which is known to lead to network hotspots. We propose a set of dynamic non-uniform resource reservation policies and dynamic routing methods. One of the routing methods is the well-known widest-shortest path method and the other is a dynamic routing method that adaptively adjusts link costs and uses a similar algorithm to shortest-path routing (e.g. Dijkstra's algorithm). We show that for both uniform and non-uniform traffic loading of some example network topologies that the combination of the proposed resource reservation policies and dynamic routing can lead to significant reduction in the connection blocking ratio in all loading conditions except for excessively high loads..**

***Index Terms*— Call admission control, deterministic delay bounds, quality of service routing, rate-controlled earliest-deadline first scheduling, resource reservation.**

***EDICS*— 5-QOSV Quality of Service, 5-BEEP Multimedia Traffic Management.**

I. INTRODUCTION

High-speed multi-services packet networks have created opportunities for many applications such as real-time distributed computation/simulation, scientific visualization, digital continuous media, and high-precision medical imaging. Such applications demand quality-of-service (QoS) guarantees on message delivery in terms of bounded delay or loss. These QoS guarantees are usually negotiated at the connection establishment time and (if connection is admitted to the network) must be adhered to by the network in the presence of background non-real-time traffic.

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. For applications requesting deterministic bounds on end-to-end packet delay, many scheduling disciplines have been proposed in the literature to implement such mechanism (see, [15], [16], [22]). Each scheduling discipline requires algorithms for performing call admission control (CAC) and resource reservation. These mechanisms have been originally proposed for ATM networks which are designed to support application-level guarantees. However, if IP were to be the universal protocol of choice for multi-services packet networks, IP-based packet networks would be required to support such hard bounds on QoS for some services. We envision that in a large public network that with the IP differentiated services (DiffServ) framework [4], an IP-based network could provide such bounds to aggregated streams mapped to the DiffServ expedited forwarding (EF) class. Moreover, within a private IP-based network or at the edge of public networks, the integrated services (IntServ) framework [4] along with RSVP signalling [6] can still be used to support application-level QoS.

A. *Contributions of the Paper*

This paper addresses the generic problem of resource reservation and routing for connections requiring a deterministic bound on end-to-end delay where RC-EDF service discipline (as proposed in [15]) is used. The aforementioned connections could be an aggregated flow (as in a DiffServ EF flow or ATM virtual path) or an application-level flow or set of flows (as in an ATM virtual channel or RSVP-based IP network). For sake of generality, we have intentionally used model-based traffic (rather than trace-based traffic) applied to some general network topologies to be able to reach general conclusions not specific to particular applications. The paper focuses on how to select an appropriate path for an incoming connection and how to map the end-to-end delay requirement of a connection into a local resource requirement to be reserved at each node along the connection's path.

A commonly used resource reservation policy is to equally divide the end-to-end delay requirement among the nodes serving the connection so that each node reserves the same

Manuscript received November, 2002, revised May 2003, November 2003. Accepted December 2003.

K. Elsayed is with the Department of Electronics and Communications Engineering, Cairo University, Egypt 12613 (phone: +20-10-525-2907; fax: +20-2-572-3486; e-mail: khaled@ieee.org).

amount of resources regardless of its capacity (i.e. the speed of its outgoing inks or its loading state). Moreover, for such services, the commonly used shortest-path routing is usually not a good choice for meeting the requested quality of service. We propose a number of non-even reservation policies that take into consideration the selected nodes link capacities and loading state along the path. We also propose using widest-shortest path or adaptive dynamic routing to enhance the performance of delay-bounded service by reducing the connection blocking rate (this occur by appropriately selecting a candidate path for a new connection over which it is more likely to be able to setup the connection). Since RC-EDF is not a rate-based scheduling discipline, most of the work done for QOS-routing for rate-based disciplines [20], [7] such as packet-by-packet generalized processor sharing (PGPS) does not directly apply. We would like to stress here the fact that our focus in this paper is to provide means for realistic resource reservation policies and routing algorithms. More theoretical treatment of the resource reservation policies for RC-EDF could be found in a related work [2] in which the non-uniform RC-EDF resource reservation policies are first proposed and their performance is assessed. In this work, we show that the proposed resource reservation and routing algorithms increase the probability of accepting new connections significantly when compared with the uniform resource reservation and shortest-path routing for both uniform and non-uniform traffic under light to moderate network loading conditions.

The rest of this paper is organized as follows: In section 2, we provide overview of related work. Section 3 presents the system model and details of the EDF scheduling discipline, CAC operation, and route selection methods. In section 3, we also provide overview of the proposed framework for deterministic end-to-end delay service provisioning including resource reservation policies for RC-EDF schedulers, dynamic routing and CAC. Section 4 presents simulation results illustrating the realized gains achievable by the methods. Section 5 concludes the paper.

II. RELATED WORK

There have been several attempts to handle the problem of resource reservation and call admission control for packet networks providing deterministic service. In an early work [24], it is shown that local deterministic delay bounds can be guaranteed over a link for bursty traffic even when the sum of the peak rates of all the connections is greater than the link speed. This allows a multi-fold increase in the number of admitted connections when the traffic is bursty. The results can be efficiently extended from a single switch to a network of arbitrary topology by using rate-controlled service disciplines at the switches.

In [21], development of a nodal metric defined as the "relative gain ratio" that predicts the relative performance of QoS allocation policies in a network setting is presented. Computation of the relative gain ratio and direct evaluation of allocation policy performance for two simple network models are done. It is found that applications which tolerate large packet loss or alternate QOS metrics, QOS allocation policies differ significantly in their performance.

Ref. [11] 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.

Ref. [19] 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 that 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.

This work is an extension of the work presented by the author and other collaborators in [2]. In [2], the non-uniform RC-EDF resource reservation policies are proposed and their performance is evaluated for simple network models with fixed or random routing. A mathematical expression is derived to express the achievable gain for a policy called the optimal static policy over the uniform allocation policy in a simple single path network of $N \geq 1$ nodes. It is also shown that in general, there is much to be gained from using non-uniform reservation policies. Also, Ref. [9] studies the problem of resource allocation and CAC in a network with PGPS scheduling. Various non-uniform resource allocation policies were handled and their qualitative performance compared.

Important contributions have also been made in the area of statistical bounds for end-to-end delay services. In [17], a scheme called coordinated multihop scheduling (CMS) is proposed. This work is based on the thesis that when packets incur low delays at upstream nodes, downstream nodes can reduce their priority and schedule other packets first. Similarly, if packets are subject to heavy congestion at upstream nodes, downstream nodes can provide higher priority to the suffering packets. CMS is shown to provide tight statistical bounds for end-to-end delay and potential performance advantage over the PGPS, EDF, and FIFO disciplines. There is some resemblance between the CMS scheme and our proposed reservation policies since both works try to exploit inter-node dependencies and/or heterogeneity in node capacities to enhance the performance. However, there are fundamental differences between CMS and our proposed schemes. Firstly, our work is focused on deterministic guarantees versus statistical bounds in CMS. Secondly, we use per-connection state in the network (i.e. all packets belonging to the same connection are treated similarly) while in CMS each packet contains its own state and priority information. Thirdly, we also include a routing perspective where we investigate the effect of various routing strategies on the performance. Work related to CMS include, coordinated-EDF [1], generalized-EDF [8], and core-jitter virtual clock [23]. It is shown in [17] that these schemes can be characterized as CMS disciplines.

An overview of quality-of-service routing is provided in [7]. Several QOS routing algorithms were evaluated by [20]. Due to the nature of the RC-EDF discipline, most of the well-known QoS routing algorithms for delay-bounded service cannot be applied directly and constrained shortest-path routing should be used. However, constrained shortest-path routing is known to be an NP-hard problem [7]. We use routing methods which add little overhead to the standard shortest-path routing while achieving orders of magnitude improvement in performance (as will be shown in Section 4).

III. THE PROPOSED FRAMEWORK FOR RESOURCE RESERVATION POLICIES, ROUTE SELECTION AND CONNECTION ADMISSION

We consider a network with arbitrary topology represented by a graph $G=(S, E)$ where S is the set of network nodes 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. However, packets from delay-guaranteed connections can preempt other packets. Under these assumptions, we can isolate the delay-guaranteed connections from other connections and analyse them separately.

Each connection j is shaped by a token bucket traffic shaper and its traffic is specified by either the (σ_j, ρ_j) descriptor or the (σ_j, ρ_j, c_j) descriptor, where σ_j is the maximum burst size (in the appropriate units of bits or cells in ATM case), ρ_j is the sustained rate (in bits/sec or cells/sec), and c_j is the peak rate (in bits/sec or cells/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 $c \to \infty$, $a \to 0$, but $ac \to \sigma$.

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

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

Maintaining the delay guarantee made to connection j is equivalent to having all the packets belonging to it transmitted completely before their assigned deadlines. Consequently, all packets from connection j do not get delayed beyond the delay bound reserved for connection j 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 a condition that must hold true for all work-conserving disciplines in general. If the condition does not hold true, a scheduler with finite buffers will always overflow and drop packets. For a scheduler k with capacity C^k (which is the data rate in bits/sec of the link following the scheduler that is reserved for delay-guaranteed connections), the stability condition is:

$$\sum_{j=1}^{N} \rho_j \le C^k \qquad (1)$$

where ρ_j is the average rate (in bits per second) of the traffic source of connection j and N is the number of connections that are being served by scheduler k.

2- The schedulability condition: The schedulability condition guarantees that a scheduler will not make violations and will, therefore, honour the QoS commitment made during the connection set-up phase. The form of the schedulability condition differs, in general, for each scheduling discipline. If the schedulability condition is sufficient but not necessary, then there is a possibility of under-utilizing the scheduler since a violation of the sufficient conditions does not necessarily mean that a given set of connections is not schedulable. On the other hand, a schedulability condition that is both necessary and sufficient guarantees that there is no under-utilization of the scheduler's resources.

A. *The EDF schedulability conditions for token-bucket shaped traffic*

Here, we use the results presented in [12] which applies Theorem 1 in [18] to deduce the schedulability conditions for token-bucket traffic models. The analysis in [12] 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, [12] defines the function $F(t)$ as:

$$F(t) = Ct - \sum_{j \in N} A_j^*(t - d_j) \qquad (2)$$

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 [12], it is shown that the *necessary and sufficient* schedulability condition of an EDF scheduler is equivalent to verifying that:

$$F(t) \geq 0 \qquad \forall\, t \geq 0 \qquad (3)$$

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

B. Details of the resource reservation, dynamic routing, and call admission schemes

The scheme we propose provides a general framework for CAC, path selection and resource reservation. The path selection affects the overall performance of the CAC process. A routing protocol that chooses a longer path or a highly saturated path for an incoming connection can cause the likelihood of connection blocking. Since it is very complex to combine routing and CAC, our assumption is that a path is first selected by a routing algorithm and then the CAC is applied on the selected path. Herein, we describe the assumed CAC operation when accepting a new connection and then discuss the proposed resource reservation policies applied within the CAC and the proposed path selection algorithms.

Connection j arrives at node n with traffic descriptor (σ_j, ρ_j, c_j) or (σ_j, ρ_j), required end-to-end delay bound D_j, and destination $m \neq n$. The routing algorithm selects a path P_j with K_j schedulers (or links) between n and m. The CAC process then decides in accordance with a certain reservation policy in place whether connection j can be admitted or not.

For bounded delay service, the application of the CAC algorithm on a connection’s path requires the computation of the minimum delay that each scheduler along the path of the connection can guarantee to this new connection. This allows the CAC algorithm to determine the minimum achievable end-to-end delay bound for this connection, and thus to determine if the network can guarantee the requested delay bound or not. Define d_j^{i*} as the minimum delay that a scheduler i that belongs to path P_j can guarantee to connection j and $D_j^* = \sum_{i=1}^{K_j} d_j^{i*}$ as the minimum achievable end-to-end delay for connection j over path P_j. (The method for calculating the minimum delay that a scheduler can guarantee to a connection specified by the (σ, ρ , c), and (σ, ρ) traffic models is given in [12].)

The CAC operation proposed in [11], [12] 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 j, sends a SETUP message to the called party, including the connection’s (4) traffic characteristics (σ_j, ρ_j, c_j), and the required end-to-end delay bound D_j. This message travels over the K_j schedulers belonging to the path P_j selected for the connection by the routing algorithm in use. At each scheduler i on P_j, the minimum delay that the RC-EDF scheduler i can guarantee to connection j, d_j^{i*}, is computed and added to d_j^*, which is the cumulative delay sum included in the setup message. If at some scheduler, the cumulative delay d_j^* exceeds the required delay bound D_j, then the connection cannot be accepted and a RELEASE message is returned to the calling party. Otherwise, the setup message reaches the last scheduler which checks if $D_j \geq D_j^*$. If the condition is true, the connection is accepted, and a CONNECT message is then returned on the same path to the calling party, reserving a delay bound $d_j^k \geq d_j^{k*}$ to connection j at each scheduler k such that $\sum_{i=1}^{K_j} d_j^i \leq D_j$. The values of d_j^k are chosen according to one of the delay reservation policy proposed in section 1).

We assume that only the last scheduler checks the validity of condition ($D_j \geq D_j^*$) and makes the irreversible decision of accepting or rejecting the connection.

1) The proposed RC-EDF resource reservation policies

The most commonly used policy is to allocate equal delay deadlines to all the schedulers in the selected path of the connection. This policy is called the EVEN policy and it is a static policy that does not take into account the network loading state, link capacities, or differences in schedulers capacities.

In [2], it was shown that dynamic policies outperform this simple policy in many network configurations. However, the configurations used were confined to small networks, or (a possibly large) single path network with fixed or random routing. The dynamic 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. Using the previous definitions of D_j^* and K_j of a generic connection, we define the excess end-to-end delay as follows:

$$\overline{D}_j = D_j - D_j^* \tag{5}$$

We now show how the local delay bound is allocated in each of the proposed policies.

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

$$d_j^i = \frac{D_j}{K_j}, \ \forall i \in \{1, \cdots, K_j\} \tag{6}$$

Even distribution of excess delay (DYNEVEN): This policy is the one suggested in [10]. The delay bound formula is given by:

$$d_j^i = d_j^{i^*} + \frac{\overline{D}_j}{K_j}, \forall i \in \{1, \cdots, K_j\} \tag{7}$$

Capacity proportional distribution of excess delay (DYNCP): 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:

$$d_j^i = d_j^{i^*} + \frac{\overline{D}_f / C^i}{\sum_{n=1}^{K_j} 1/C^n}, \forall i \in \{1, \cdots, K_j\} \tag{8}$$

Remaining-delay proportional distribution of excess delay (DYNRDP): 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:

$$d_j^i = d_j^{i^*} + \frac{\overline{D}_j}{D_j^*} d_j^{i^*}, \forall i \in \{1, \cdots, K_j\} \tag{9}$$

$$\Rightarrow d_j^i = d_j^{i^*} \frac{D_j}{D_j^*}, \forall i \in \{1, \cdots, K_j\} \tag{10}$$

2) *Dynamic routing schemes*

The most commonly used routing protocol in data networks is shortest-path routing. Except for the case of lightly loaded networks, shortest-path routing is in general not suitable for QoS oriented multi-service networks. Moreover, it suffers from leading to hotspots and of being susceptible to inefficient load balancing in a given network [3],[14]. The problem of constrained routing (for example route selection subject to a maximum end-to-end delay as the problem addressed here) is NP-hard. Several heuristic algorithms exist for reducing the complexity of the route selection. However, to the best of our knowledge nothing specific has been devised to handle the case of networks with RC-EDF schedulers. Most of the work proposed for this area assumed the use of variants of the PGPS scheduling discipline (see [20] for example).

In this work, we evaluate the performance of the proposed resource reservation policies using the following routing algorithms.

Shortest Path First Routing (SP): The first routing method is the standard shortest-path first routing method, which is extensively used in protocols such as OSPF and RIP. We fix the link costs to 1, in which case shortest-path is equivalent to shortest-hop routing. We use the standard Dijkstra algorithm to implement the SP algorithm.

Widest-Shortest Path Routing (WSP): Widest-shortest path routing [24] is a popular method for quality-of-service routing. In WSP, when two paths with equal costs to a certain destination are found, the one with the highest available capacity is chosen. We modify the Dijkstra algorithm to keep track of the available capacity on each link in order to implement the WSP algorithm. In our version of WSP, the available capacity on a link is equal to the link capacity minus the sum of average rates of connections passing through the link. This information is readily kept on each link for the necessary book-keeping needed by the CAC process.

Dynamic Routing (DR): In this method, we use the standard Dijkstra algorithm for shortest-path routing while setting the link cost adaptively. Let the utilization of link *e* be given by $U^e = \sum_{j=1}^{N_e} \rho_j / C^e$, where N_e is number of connections passing though link *e,* we set the cost of link *e* to be $1/(1-U^e)$. Our early experimentation with this choice for link costs has proven to provide excellent load balancing capabilities and overcomes hot-spot problems that can be created with standard shortest-path routing. Also, using a similar piece-wise linear cost function has proven as a good choice to do traffic engineering in IP networks with traditional interior gateway routing protocols (OSPF) as reported in [14].

Note that for both WSP and DR, the update in the available capacity (for WSP) or link cost (in DR) happens only after a new connection is accepted to the network and when an existing connection is terminated. The two algorithms would need comparable signalling overhead. However, DR is based on executing the well-known Dijkstra algorithm, while WSP needs more modifications.

IV. PERFORMANCE EVALUATION

Comparing the performance of the proposed reservation policies and routing schemes requires the definition of a performance metric against which the different policies may be compared. A suitable performance metric would be 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 to compare the different policies for this study. We use the *simpack* [13] package for discrete event simulation.

A. Simulation details

The overall connection arrival process to the network is generated according to a Poisson process with an average arrival rate of λ connections/sec. Connection duration is assumed to be exponentially distributed with a mean $1/\mu$ seconds. The value $\rho=\lambda/\mu$ characterizes the offered traffic load (in Erlangs) at the network under consideration. We fix μ to be one unit. The source and destination nodes for a newly arriving connection is chosen according to whether we have uniform or non-uniform traffic as will be discussed later. For each simulation, we generate a total of between 100,000 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-4). The blocking probability is computed as the number of blocked connections divided by the number of generated connections.

The parameters of the simulation, most notably link capacities, end-to-end delays, and source descriptors; are carefully chosen to meet the following objectives where possible: 1- the delay requirements are as close as possible to the delay requirements of interactive voice or video connections over multi-service packet networks (e.g. ATM); 2- traffic parameters are similar to the traffic characteristics of standard packetized voice and video encoders; and 3- the schedulability condition of a scheduler, rather than the stability condition, is the one that puts the limit on the maximum number of connections. This is done because the performance of several resource reservation policies can only be differentiated when the stability condition does not limit the number of connections.

Our simulation supports both the (σ, ρ) and the (σ, ρ, c) models. However, the paper only presents the results of the (σ, ρ) traffic model. Realistic values of traffic characteristics are used and are selected as follows [12]:

- 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]

The 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 *Fig. 1*). We evaluate the performance for two link configurations. The first configuration is when all links have the same capacity of 34 Mbps (equal-capacity links). (It should be noted that for the equal-capacity links case, the DYNEVEN and DYNCP policies are equivalent.) In the second configuration, link capacities are randomly chosen such that overall capacity in the network is identical to the equal-capacity links case.

For the case of uniform (symmetric) traffic, the newly arriving connections are divided with equal probability among all possible source destination pairs in a given topology. For each topology and link configuration, we obtain the following graphs in the uniform traffic case: 1- a graph showing the blocking probability for the various policies (EVEN, DYNEVEN, DYNCP, DYNRDP) with shortest-path routing; 2- a graph showing the blocking probabilities for the best performing policy for each of the three used routing algorithms (SP, WSP, and DR). The best performing policy is found as described below; and 3- a graph showing the reduction in the blocking probability achieved by using WSP or DR with the corresponding best policy for each as compared to that obtained when SP is used in conjunction with best performing policy (gain1) and as compared when SP routing is used in conjunction with the EVEN policy (gain2).

The best performing policy is obtained as follows. For each routing method, we find the minimum achievable blocking probability among the four policies for each value of the traffic load (minimum envelope). We then evaluate a policy distance metric, which is the root mean square of the differences between the blocking probability values of the policy and the minimum envelope. The policy with the smallest distance metric is the best performing policy. (In many cases, one policy completely defines the minimum envelope, and its metric is thus zero.)

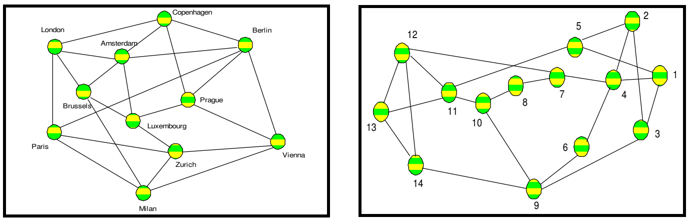

(a) The COST239 Network Topology (b) The NSFNET Network Topology

Fig. 1. Network topologies of the COST239 and NSFNET networks.

B. *Results for the COST239 network*

We change the input traffic from 2 upto 2048 Erlangs and obtain the blocking probability for the different resource reservation polices. The results for the equal-capacity and random link capacity allocation are shown in Fig. 2 (a, b, c) and (d, e, f) respectively. For both the equal-capacity and random links case with SP routing, we note that the dynamic policies outperform the EVEN policy especially for low values of traffic load. This is also true for WSP and DR as well. The best performing policy is always the DYNCP policy for all routing methods and link configurations. This is a good finding since link capacities are usually constant and in contrast to the DYNRDP policy, the signalling overhead needed for DYNCP would be much lower. By comparing the performance of this best-performing policy under the three routing methods, we see that DR and WSP routing offer significant reduction in blocking probability in moderate and high loading. Moreover, for the case of random link capacities, this reduction can be upto 50-60% which illustrates the need for this type of routing when there is unbalance in the network link capacities. If we inspect the curves marked DR-gain2 and WSP-gain2 in Fig. 2(c) and (f), we see a remarkable improvement for DYNCP/DR and DYNCP/WSP over what could be achieved by using EVEN/SP (upto 78% reduction in the random links case).

Comparing WSP and DR, we find that WSP seems to provide more resilient performance specially for the random links case with low values of traffic load. In medium to high traffic load conditions and in the case of equal-capacity links, DR usually offers superior performance. Our choice seems to prefer DR routing since its complexity is less the WSP routing.

C. *Results for the NSF network*

The NSF network topology has a lower connectivity than the COST239 network (as measured by number of links/node). Similar to what we have done in the COST239 network, we obtain the results for the NSF topology with equal-capacity and random links case in Fig. 4(a, b, c) and (d, e, f) respectively. Similar behaviour is observed here as well (with higher blocking rates in general). The DYNCP is still the best performing policy for all routing methods and link configurations. By inspecting the results for the reduction in blocking probability, we see that only for the case of random links with low loading, that SP outperforms DR. However, WSP always achieves some reduction in blocking probability. In comparison with the EVEN/SP combination, we see a remarkable improvement by using DYNCP/DR or DYNCP/WSP.

We highlight here that the choice between DR and WSP is a tradeoff between complexity and performance. DR has the advantage of being based on implementing a simple dynamic cost function while WSP is more complex to implement. Our choice seems to prefer DR routing.

D. *Results for non-uniform traffic loading*

We study the COST239 network topology under non-uniform traffic loading. To generate non-uniform traffic, we first define the traffic non-uniformity factor as follows. We identify a set of *hot* source destination pairs. These hot pairs are assigned a portion α of the overall traffic. Thus, by varying the number of hot pairs and the portion of assigned traffic, we can create non-uniform traffic in the network. In the COST239 networks we assume that there are 11 hot source destination pairs of the overall 110 pairs and we vary α from 0 to 0.55. We study the effect of varying α for two values of the overall arrival rate: λ=8 connections/sec and λ=512 connections /sec for equal-capacity and random link capacities.

The full set of results shows again (as in the uniform traffic case) that the DYNCP policy has superior performance over all other policies. We report the blocking probability as function of α for the two values of the arrival rate at **Fig. 4** for the EVEN policy with shortest-path routing and the DYNCP policy with shortest-path, WSP, and DR routing. It is clear that, as in the case of uniform traffic, we have significant gain from using the dynamic resource reservation policies such as DYNCP with WSP and DR routing. In general as the non-uniformity in traffic increases while fixing the arrival rate, the more blocking occurs and the more apparent the need for the proposed resource reservation policies and routing methods to reduce the blocking rates.

V. CONCLUSIONS

We proposed and evaluated the performance of dynamic resource reservation protocols and dynamic routing for multi-service packet networks employing RC-EDF scheduling to provide deterministic end-to-end delay bounds. The performance evaluation was done on general network topologies and realistic traffic parameters covering packetized voice and video. Our results show that under uniform and non-uniform traffic, the DYNCP is the best performing policy for all cases considered. This is a good finding since link capacities do not change frequently and thus the signalling overhead needed for DYNCP would be much lower than the DYNRDP policy. Moreover, widest-shortest path and dynamic routing based on link cost adaptation provide significant reduction in blocking probabilities. The choice of DR or WSP is a tradeoff between complexity and performance. DR has the advantage of using a simple cost function while WSP is more complex to implement and more robust.

It remains to be explored whether the same conclusions hold true and what performance levels could be achieved by rate-based disciplines such as PGPS. Also, a possible area of research is to extend this framework to the case of statistical delay guarantees.

REFERENCES

[1] M. Andrews. Probabilistic end-to-end Delay Bounds for Earliest Deadline First Scheduling. Proceedings of IEEE INFOCOM 2000, March 2000.

[2] A. Ayad, M. El-Hadidi, K. Elsayed, Resource Division Policies for EDF Scheduling in ATM Networks, Proc. of the IEEE ISCC'01, pp. 249-254, July 2001.

[3] D. Bertsekas and R. Gallager, Data Networks, Prentice-Hall, 2nd edition, 1992.

[4] S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang, and W. Weiss, An Architecture for Differentiated Services, IETF RFC 2475, Dec. 1998.

[5] R. Braden, D. Clark, and S. Shenker, Integrated Services in the Internet Architecture: an Overview, Internet RFC 1633, June 1994.

[6] R. Braden, L. Zhang, S. Berson, S. Herzog, S. Jamin, Resource ReSerVation Protocol (RSVP) Version 1 Functional Specification, Internet RFC 2205, Sep. 1997.

[7] S. Chen and K. Nahrstedt, An Overview of Quality of Service Routing for Next-Generation High-Speed Networks: Problems and Solutions, IEEE Network, Vol. 12, pp. 64 -79, Nov./Dec. 1998.

[8] T. Chen, J. Walrand, and D. Messerschmitt. Dynamic Priority Protocols for Packet Voice, IEEE Journal on Selected Areas in Communications, Vol. 7(5):632–643, June 1989.

[9] K. Elsayed, A. Ayad, and M. El-Hadidi, Performance Evaluation of Resource Reservation and Call Admission Policies for Deterministic Services in PGPS-based Packet Networks, Elsevier Computer Communications, Vol. 25, pp. 1513-26, Nov. 2002.

[10] D. Ferrari and D. C. Verma, A Scheme for Real-Time Channel Establishment in Wide-Area Networks, IEEE Journal on Selected Areas in Communications, Vol. 8, No. 3, pp. 368-379, 1990.

[11] V. Firoiu and D. Towsley, Call Admission and Resource Reservation for Multicast Sessions, Proceedings of IEEE INFOCOM'96, pp. 94-101, March 1996.

[12] V. Firoiu, J. Kurose, and D. Towsley, Efficient Admission Control of Piecewise Linear Traffic Envelopes at EDF schedulers, IEEE Trans. on Networking, Vol. 6, No. 5, pp. 558-570, Oct. 1998.

[13] P. Fishwick, The Simpack ToolKit, available from http://www.cise.ufl.edu/~fishwick/simpack.html.

[14] B. Fortz and M. Thorup, Optimizing OSPF/IS-IS Weights in a Changing World, IEEE JSAC, Vol. 20(4), pp. 756-767, .May 2002.

[15] L. Georgiadis, R. Guerin, V. Peris, and K. Sivarajan, Efficient Network QoS Provisioning Based on Per Node Traffic Shaping, IEEE Trans. on Networking, Vol. 4, pp. 482 -501, Aug. 1996.

[16] J. Golestani, A Framing Strategy for Congestion Management, IEEE Journal on Selected Areas in Communications, Vol. 9, No. 7, Sep. 1991.

[17] C. Li and E. Knightly, Coordinated Multihop Scheduling: A Framework for End-to-End Services, IEEE/ACM Transactions on Networking, Vol. 10(6), pp. 776-789, December 2002

[18] J. Liebeherr, D. E. Werge, and D. Ferrari, Exact Admission Control for Networks with a Bounded Delay Service, IEEE Transactions on Networking, Vol. 4 No. 6, pp. 885-901, 1996.

[19] D. H. Lorenz and A. Orda, Optimal Partitioning of QoS Requirements on Unicast Paths and Multicast Trees, IEEE Trans. on Networking, Vol.10(1), pp. 102-114, 2002.

[20] Q. Ma and P. Steenkiste, Quality-of-Service Routing with Performance Guarantees. 4th International IFIP Workshop on Quality of Service, May 1997.

[21] 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.

[22] A. K. Parekh and R. G. Gallager, A Generalized Processor Sharing Approach to Connection Control in Integrated Services Networks: The Multiple Node Case, IEEE Trans. on Networking, vol. 2, pp. 137-150, 1994.

[23] I. Stoica and H. Zhang. Providing Guaranteed Services Without per Flow Management, Proceedings of ACM SIGCOMM '99, August 1999.

[24] Z. Wang and J. Crowcroft, QoS Routing for Supporting Multimedia Applications, IEEE JSAC, Vol. 14, pp. 1228-1235, September 1996.

[25] H. Zhang and D. Ferrari, Improving Utilization for Deterministic Service in Multimedia Communication, Proceedings of the International Conference on Multimedia Computing and Systems, pp. 295–304, 1994.

Khaled M. F. Elsayed (S'90–M'95-SM'02 received his B.Sc. (honors) in electrical engineering and M.Sc. in engineering mathematics from Cairo University in 1987 and 1990 respectively. He received his Ph.D. in computer science and computer engineering from North Carolina State University in 1995. He is now an Associate Professor in Cairo University, Egypt and is an independent telecommunications consultant. Between 1995 and 1997, he was a member of scientific staff with Nortel Wireless Systems Engineering in Richardson, TX. Dr. Elsayed was the editor for the Internet technology series of the IEEE Communications Magazine from 1998 until 2002. He has served on technical program committees for several IEEE, IFIP, and ITC conferences. He was the technical co-chair for IFIP MWCN 2003 conference in Singapore. He also served as an expert evaluator for the European Commission FP5 and FP6 programmes.

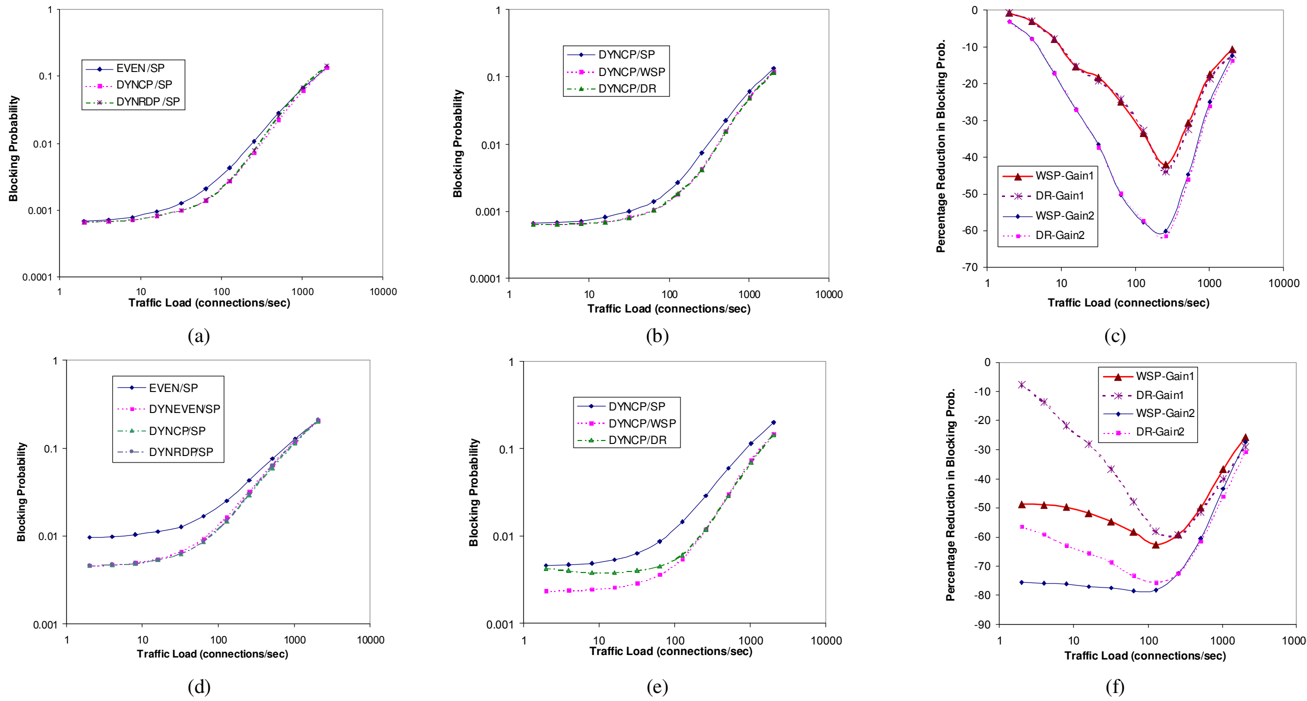

(a) (b) (c)

(d) (e) (f)

Fig. 2. Results for the COST239 network with uniform traffic and equal-capacity links (a,b,c) and with uniform traffic and random links (d,e,f).

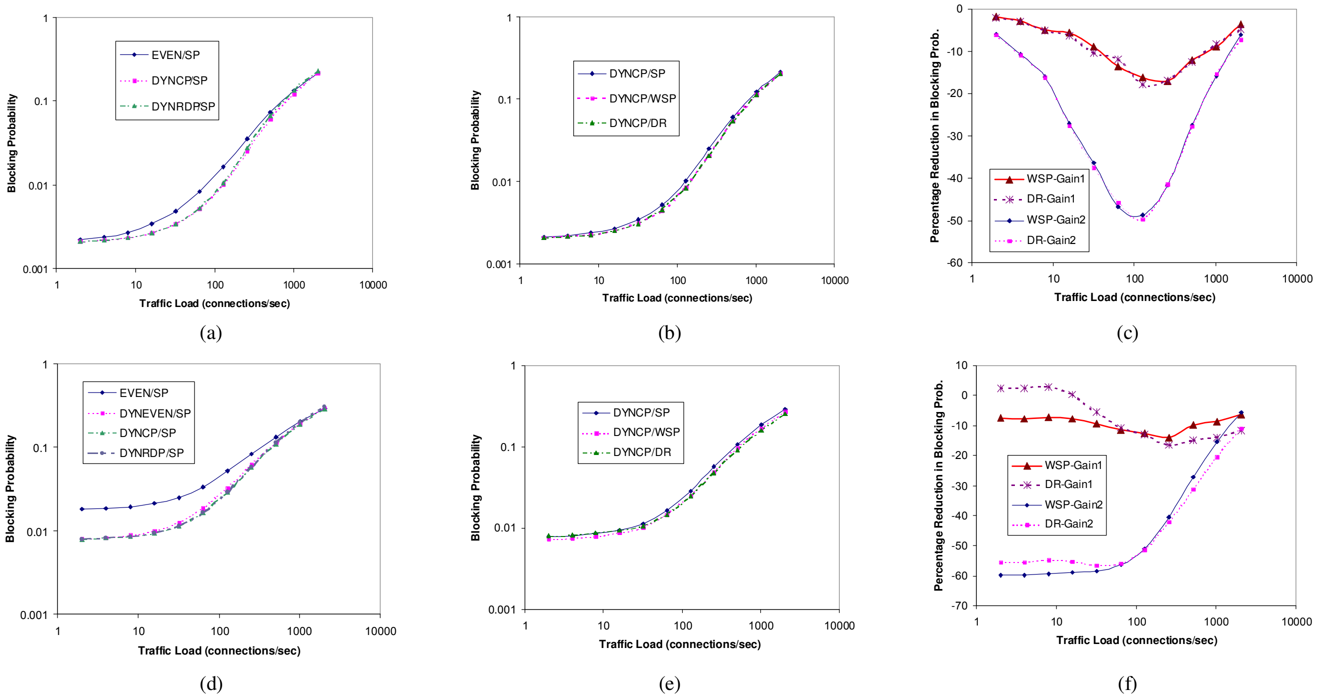

(a) (b) (c)

(d) (e) (f)

Fig. 3. Results for the NSFNET network with uniform traffic and equal-capacity links (a,b,c) and with uniform traffic and random links (d,e,f).

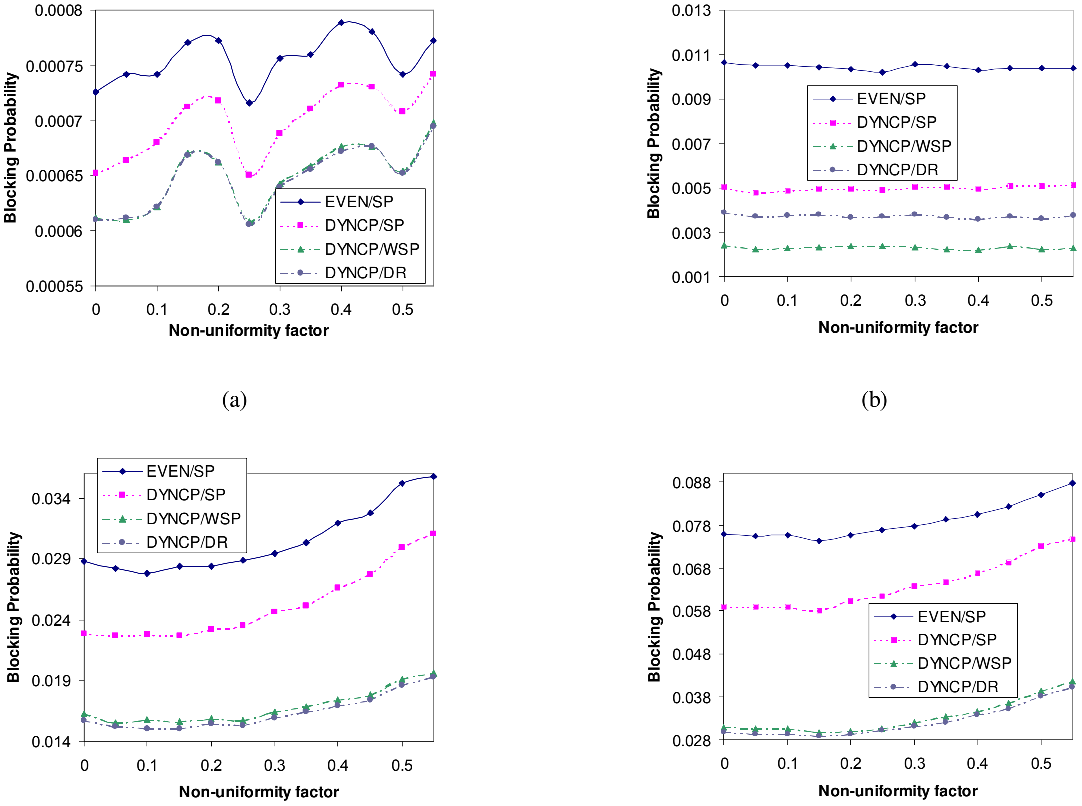

(a) (b)

(c) (d)

Fig. 4. Effect of non-uniform traffic loading on the performance of the EDF policies and the routing algorithms: COST239 network (a) arrival rate of 8 calls/sec and equal-capacity links, (b) arrival rate of 8 calls/se and random links case, (c) arrival rate of 512 calls/sec and equal-capacity links, and (d) arrival rate of 512 calls/sec and random links.