The EDF Scheduling with Multiple Loss Requirements Support

Xiaojun Hei, Danny H. K. Tsang
Dept. of Electrical & Electronic Engineering
Hongkong University of Science and Technology
Clear Water Bay, Kowloon, Hongkong
Email: {heixj, eetsang}@ee.ust.hk

Abstract

In this paper, we study the problem of QoS guarantee for different service classes on the Internet. A 2-level hierarchical scheduling framework is employed for the separation of QoS requirements such as bandwidth, delay and packet loss. The real-time applications running on the UDP transport may have significantly different loss requirements. We propose a novel modified EDF scheduling scheme (E-LOSS), in which the loss probabilities of different flows are measured on-line and used for packet scheduling when the queue length exceeds some pre-specified threshold. The simulation studies show that the proposed E-LOSS scheme can support real-time connections with different delay and loss requirements with better efficiency than the generic EDF scheduling and the modified EDF scheduling for finite buffer (EDF-FB).

Keywords: Scheduling, Earliest Deadline First, Real-time Traffic

1 Introduction

With the fast commercial deployment of the Internet in the past several years, an increasingly large number of real-time applications with diverse performance requirements in terms of different bandwidth, loss and delay demands, have been running on the Internet. For example, a CD quality music broadcast requires a very stringent delay and loss guarantee, while Internet telephony has a much more relaxed requirement. With the packet-switched technology used on the Internet, all types of traffic are multiplexed together and share the same resources in the network. These applications expect the network to be capable of supporting different Quality-of-Service (QoS) guarantees efficiently.

Due to the desirable property of minimizing the maximum packet lateness, the Earliest Deadline First (EDF) is proposed to schedule real-time traffic [1]. In the EDF scheduling scheme, packets are differentiated by priorities in terms of the deadline assigned to the packets. With a more stringent deadline, the packet is served earlier and it will be less likely to be dropped in the network. Therefore, the delay requirement of the traffic will influence its loss rate even though the traffic may have a loss requirement independent of its delay constraint.

There are two types of packet loss in the EDF scheduling: packets may be dropped due to insufficient buffer space or deadline violation when they arrive at a router. They are referred to as overflowed packets and overdue packets respectively. The overdue packets are dropped based on the assumption that they are useless even if they arrive eventually at the destination in a real-time communication. In order to support multi-loss requirements efficiently, both buffer management and packet scheduling schemes should be carefully designed.

Considering the information about buffer occupancy, the overflow probability can be reduced while an optimal delay performance for the system can still be maintained by the Earliest Deadline First scheduling scheme for Finite Buffer (EDF-FB) [9]. EDF-FB implements a per-flow buffer structure and a threshold is set for each buffer. If the queue length is greater than this threshold, the flow with the longest queue is served. The simulation study shows, however, that heterogeneous traffic with different delay requirements experiences approximately the same loss performance.

The buffer management scheme, proposed in [6] and [7], can support multiple delay and loss requirements in ATM networks with a measurement based cell discarding policy, in which the on-line measurement of cell loss ratio of each flow is maintained and when a cell arrives at a full buffer, one cell is dropped from the flow with the smallest ratio between its loss measurement and requirement. This cell drop scheme actually punishes the flow with the best normalized loss performance, with the expectation that cells belonging to other flows have greater chances to go into the buffer and hence grasp more capacity, so that their loss rates can be decreased in the future. However, the link capacity may not necessarily be utilized by other flows if there are no further arrivals from those flows. Instead of using the on-line loss measurement as the base of packet discarding, we propose that, if the flow is expected to violate its loss requirement, the scheduler can serve this flow first.

In this paper, a novel modified EDF scheduling scheme (E-LOSS) is proposed in order to support multiple delay and loss requirements. With the expectation for the current Internet to be able to deliver packets with a diverse set of QoS guarantees while achieving a high bandwidth utilization, the packet loss probability is also measured on-line and used for packet scheduling when the queue length is greater than some pre-specified threshold.

Because of the tight latency bound in real-time applications, it is suggested that they had better not use a reliable transport like TCP because retransmission of lost packets may cause packets to arrive too late to be useful [8]. Instead, the unreliable transport, such as UDP, is recommended. Therefore, we assume that the real-time traffic is only carried by the UDP transport and hence E-LOSS is deployed for multiple UDP datagrams which compete for a guaranteed link capacity.

The paper is structured as follows. In Section 2, we present the 2-level hierarchical scheduling framework in order to support multi-class traffic running on the Internet at the same time. Section 3 describes how the modified EDF scheduler (E-LOSS) works. The simulation based performance evaluation is then carried out in Section 4. Finally, we conclude the paper and present the future work in Section 5.

2 Hierarchical Scheduling Framework

The hierarchical scheduling aims to meet the goals of sharing link capacity and providing differentiated service, such as real-time service, best-effort service, and others [3]. We propose the hierarchical scheduling framework to explore the property for the EDF scheduler to minimize the maximum packet latency [10], where the higher level packet-by-packet General Processor Sharing (PGPS) scheduler tries to fairly allocate bandwidth among classes and either PGPS or EDF scheduler at the lower level is chosen to maintain the particular throughput or delay QoS guarantee (see Figure 1). Due to the scalability concern, the scheduling framework tries to provide service with a class-based QoS. The higher level scheduler, called the link-sharing scheduler, allocates bandwidth among classes. The lower level scheduler tries to serve a traffic class with its allocated share of bandwidth. The EDF scheduler is proposed to schedule real-time traffic to meet its stringent time constraint.

The traditional best-effort traffic will still run on the future Internet and there may be more service classes emerging. We are considering how to provide service for the best-effort and real-time traffic in the 2-level hierarchical scheduling framework. In the best-effort branch, the CHOKe buffer management algorithm can also be implemented for a more fair bandwidth allocation [2]. In order to differentiate real-time traffic, the scheduler's structure should be modified into the form shown in Figure 2. A per-class or even a per-connection buffer structure can be built up to store real-time packets and work with E-LOSS.

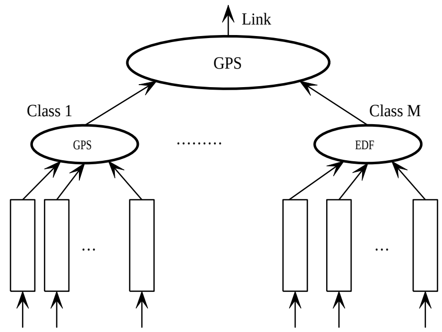

Figure 1: Hierarchical scheduling framework.

If one more traffic class appears, it should be easy to add one more branch in the hierarchical scheduler tree. To simplify the problem, we have focused our attention only on the E-LOSS branch in the tree. The real-time traffic class can be guaranteed with the worst case of minimum bandwidth by the upper PGPS scheduler. If the higher level PGPS scheduler allocates its excess bandwidth to the real-time traffic, its QoS provision will be better.

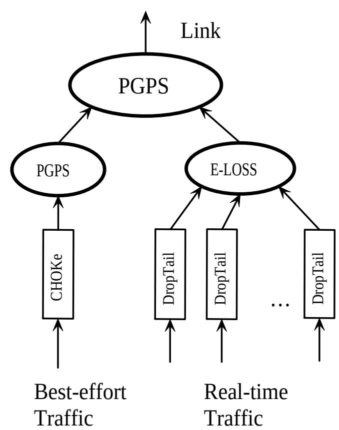

Figure 2: Hierarchical scheduling structure with diverse performance requirements support.

3 Scheduling Algorithm Design

Consider the following scenario. There are N UDP flows accessing an Internet router. The QoS parameters are specified as the delay budget (d_i), which is the total delay bound permitted for a packet to traverse the network to its destination, and the loss probability (ϵ_i), which can be measured by the ratio between the amount of loss packets and the total amount of transmitted packets. A packet is discarded when it is delayed over the pre-specified deadline or the buffer is full upon its arrival.

E-LOSS is designed to maintain a good delay performance while meeting the packet loss requirement. The block diagram of the E-LOSS scheduling scheme can be depicted in Figure 3, in which LMS represents Loss Measurement Server. Two non-simultaneous servers, EDF and LMS, are used to determine the service order of packets in the E-LOSS scheduler. The EDF server serves the HOL packet in each buffer with the earliest deadline when it is put into operation, while LMS measures the loss probability of each flow, and serves the HOL packet in each buffer with the largest ratio between its loss probability measurement and its loss requirement when LMS is in operation. Normally, packets are served by the EDF server;

however, if the queue length of any flow exceeds its corresponding threshold, LMS will serve the flow instead.

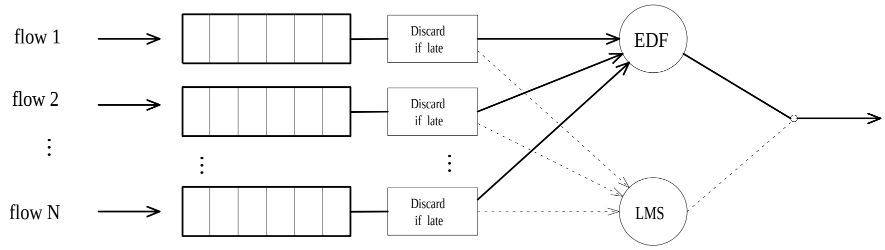

Figure 3: The E-LOSS scheduler for real-time traffic.

The detailed description of E-LOSS is illustrated in the pseudo-code below. In E-LOSS, packets from each source are queued in a separate buffer for service. In addition, a threshold is set up for queue i as the indication of buffer overflow. On arrival, each packet will be assigned a time-stamp with a value equal to the current time plus its maximum delay allowance in the queue as the deadline for the packet. The Head-Of-Line (HOL) packet with the earliest deadline is searched in a round-robin fashion with the randomly selected starting point. The random starting point in the search is proposed to obtain fair resource allocations among the flows because the simulation study shows that the flow in the front part in the search order will have more opportunities to be served earlier and have the advantage of bandwidth allocation. At the beginning of each round, all the HOL packets will be checked first. If the packet has already been delayed beyond its deadline, it will be discarded and the next packet from the same queue will be promoted to the HOL position and checked again and so on until a packet that does not violate its deadline is found. After the overdue packets are dropped, all the HOL packets in different queues are eligible for service. In normal operations, the E-LOSS scheme will act as the original EDF scheme, because the EDF server is in operation, that is, the time-stamps of all the HOL packets are compared and the packet with the earliest deadline will then be picked and served. However, when the occupancy of one or more queues is above the threshold, LMS will come into action and override the function of the EDF server. In this situation, LMS will calculate the ratio between the loss probability measurement and the loss requirement for those queues with their occupancy above their individual threshold. Then, the HOL packet from the queue with the largest ratio will be selected and served. Therefore, those queues that are in danger of loss requirement violation will get service first and be able to recover from the temporary overloading quickly.

```
begin_serve_next_packet
  //discard those packets with deadline violation
  for i =1 to N          //deadline[i]: deadline of the HOL packet in queue i
      while time > deadline[i]   //time: current time
         discard the packet;
         promote next packet to the head of the queue;
      end
  end
  //variable initialization
  queue_to_serve = 0;     //queue_to_serve: queue to be picked for service
  earliest_deadline = deadline[1];
  max_loss_ratio = 0;
  flowid = uniform(1,N); // randomly select a flow to start;
  //look for the candidate flow with the largest normalized
  //loss ratio among those queues above the threshold
```

```
  //threshold[i]: threshold value of queue i
  //normalized_loss_ratio[i]: value of the normalized loss ratio of queue i
  for i = 1 to N
      if queue_length[flowid] > threshold[flowid]
               AND normalized_loss_ratio[i] > max_loss_ratio
         queue_to_serve=flowid;
         max_loss_ratio = normalized_loss_ratio[i];
      end
      //find the queue with the earliest deadline
      //if no queue is above the threshold
      if deadline[flowid] <= earliest_deadline
         EDF_queue_to_serve = flowid;
         earliest_deadline = deadline[flowid];
      end
      flowid++;
      if flowid > N
         flowid = 1;
      end
  end
  if queue_to_serve = 0
      queue_to_serve = EDF_queue_to_serve;
  end
end_serve_next_packet
```

4 Performance Evaluation

4.1 System Topology

The simulations are carried out in Network Simulator [11]. Considering the following simulation scenario as shown in Figure 4, there are 10 sources accessing the router. All the sources are assumed to be ON/OFF sources, and the generated traffic is carried by UDP datagrams. All the packets have the same MTU size. If the last packet generated in the ON state is not long enough to fill in one full packet, the source just sends out the packet with this excess data size. The packet is dropped when it is delayed over the pre-specified deadline or the buffer is full upon its arrival.

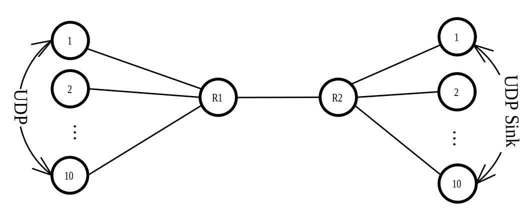

Figure 4: System topology in the simulation.

4.2 Simulation Configuration

As shown in Figure 4, the link with the capacity of 800 *Kbps* between the routers R1 and R2, is shared by 10 UDP flows. The end systems are connected to the router R1 and R2 with a 10 *Mbps* link. Due to the effect of traffic aggregation, the link between R1 and R2 becomes the bottleneck in the network. All the links in the access network (from the end system to the network) have a small propagation delay of 1 ms and the propagation delay of the link between R1 and R2 is set to 100 ms, so the delay experienced by packets is mainly caused by the queueing delay and propagation delay in the shared congested link rather than the access links. Different scheduling schemes are deployed in the congested link for performance evaluation, such as EDF, EDF-FB and E-LOSS. All packets are assumed to have the same MTU size of 100 $bytes$.

The sources of real-time traffic are modeled as ON/OFF sources with exponentially distributed ON and OFF times, which are parameterized by average burst time ($\frac{1}{\beta}$), average idle time ($\frac{1}{\alpha}$), burst rate (P) and maximum packet size (L). The burst rate is the traffic generating bit rate when the source is in the ON period. The QoS parameters specified are the maximum transmission delay (d_i) and the packet loss probability (ϵ_i), which can be negotiated with the network by signalling from the source node or can be stamped into the header of each packet before it is sent out into the network so that routers can know them by accessing the content in the packet header. The utilization (ρ) of the congested link between R1 and R2 is defined as the total average traffic input divided by the link capacity.

4.3 Effect of Threshold Setting

In this section the threshold settings of the queues are studied to see how they affect the delay and loss performance. Assume that there are 10 statistically identical ON/OFF sources ($\frac{1}{\beta} = 0.05sec$, $\frac{1}{\alpha} = 0.20sec$, $P = 320Kbps$). The maximum packet size is kept as $L = 100bytes$. The traffic has the delay requirement ($d = 0.3sec$) and the loss requirement ($\epsilon = 0.001$). The simulation lasts for 6000 sec.

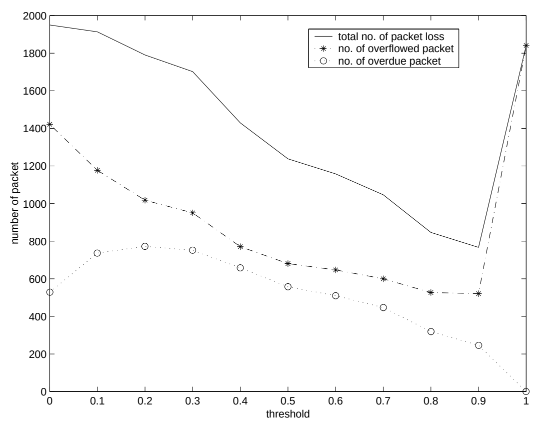

Figure 5: Packet loss v.s. threshold in E-LOSS ($\rho = 0.8$).

Figure 5 shows the effect of threshold variation upon the loss performance in the E-LOSS scheduler. A threshold of one means no buffer control. If the threshold is small, most of the time LMS will override the function of the EDF server. Since LMS only considers the total packet loss of flows based on their transmission histories as scheduling references without concerning about the actual packet deadlines, this may not lead to the optimal scheduling sequence. It is quite possible that one flow does not have the longest queue but its queue length is already over the threshold, where LMS is called into action. Although its deadline may still be a little far away, the flow is possible to have a higher loss rate due to the historical packet loss. Even though the other flows may be on the verge of buffer overflow or deadline violation, LMS will still serve this selected flow with the highest normalized loss ratio, which is not desirable in terms of fully utilizing the buffer space and minimizing the number of overdue packets. In Figure 5, there is a considerably large number of overflowed packets and overdue packets when the threshold is small. With the increase of the threshold value, the EDF server will predominantly schedule packets. However, it is still possible to build up queue length for some flow if packet deadlines in the queue are still far away; thereafter, LMS takes control again. The server switching occurs, and this will degrade the scheduling performance of the EDF server. As the threshold increases at the beginning, the number of late packets increases a little bit; however, the number of overflowed packets decreases more quickly so the total packet loss still drops. When the threshold value is increased further, more and more packets are scheduled by the EDF server. Therefore, the number of overdue packets starts to decrease and at the same time the number of overflowed packets continues to decrease. When the normalized threshold is increased to one, E-LOSS becomes the classic EDF scheduler with the per-flow buffer structure.

One interesting observation is that the number of overflowed packets decreases for the threshold between 0 and 0.8 but it increases very abruptly when the threshold is greater than 0.8 and eventually reaches 1.0. This is because when the threshold is smaller than one, the number of overflowed packets will affect the scheduling decision with the involvement of LMS. When the threshold is approaching to one, the EDF server will make the scheduling decision most of the time, which does not consider the packet dropping due to buffer overflow. Consequently, there exists a threshold setting at which the total packet loss is minimized and based on Figure 5 it had better be set at 0.8. With this threshold setting, the system achieves the best trade-off between the numbers of overflowed packets and overdue packets. In the other simulation studies, the threshold is set to be 0.8, though this optimal value is obtained with the assumption that all the sources are statistically identical. In the heterogeneous case, the optimal threshold setting may be different for different traffic classes.

4.4 Effect of Buffer Size with Homogeneous Traffic

In this section, the effect of the buffer size on packet loss is investigated under the homogeneous traffic input. Comparison studies are done between E-LOSS, EDF and EDF-FB. The system configuration in EDF and EDF-FB is set to be the same as that in E-LOSS.

Assume that all the UDP flows are statistically identical and have the same delay requirement of 300 ms and loss rate requirement of 0.01. In order to study the effects of buffer size under different traffic loading, the burst rates of the sources are changed while the other parameters are kept constant.

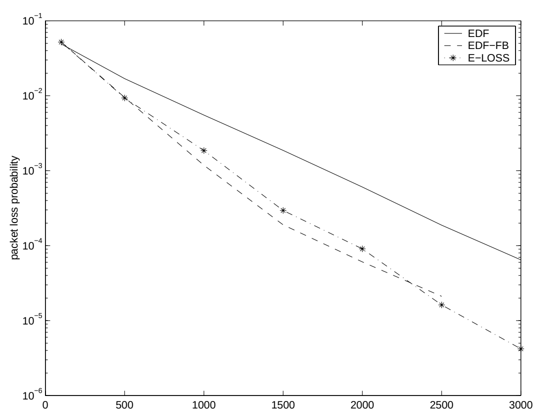

Figure 6: Loss probability v.s. buffer size under moderate loading and homogeneous traffic condition ($\rho = 0.6$, $d = 300ms$) for different schedulers.

Under moderate loading ($\rho = 0.6$), almost all packets can reach the destination before their deadlines, so the packet loss mainly arises from the packet dropping due to buffer overflow. By monitoring the packet loss, E-LOSS gives priorities to those flows which suffer more in their transmission history. Therefore, these flows can get more bandwidth in the future so that the packet loss can be minimized; hence, their overall packet loss can be decreased. In Figure 6, E-LOSS always outperforms the classic EDF scheme while it maintains a similar performance as EDF-FB. The packet loss ratio is almost an order of magnitude lower in E-LOSS than that in EDF. When the buffer size is small, the loss probability in E-LOSS is only slightly lower than that in EDF. This is because when the buffer is small, the buffer cannot hold all the incoming bursty traffic before they are served even if all the bandwidth is given to the overflowing queues. As the buffer size increases, the loss probability of E-LOSS decreases more quickly than that in EDF, so that the relative improvement of E-LOSS increases. This is because when the buffer size is large, the buffer can accommodate more packets during a burst. The overloading queues are likely to have longer queue so that LMS can be put into operation and help these overloading queues recover from temporary overloading. As a result, the loss probability in E-LOSS decreases faster than that in EDF with the increase of the buffer size.

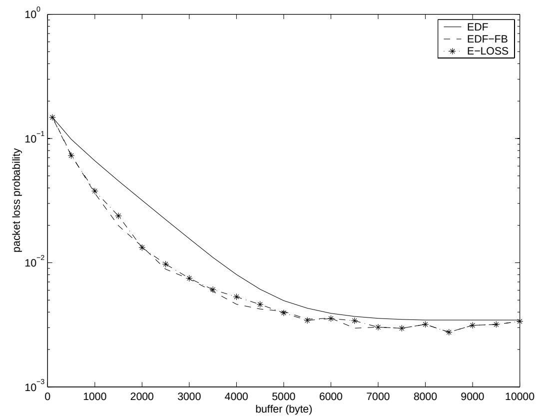

Figure 7: Loss probability v.s. buffer size under heavy loading and homogeneous traffic condition ($\rho = 0.8$, $d = 300ms$) for different schedulers.

The simulation results for heavy loading ($\rho = 0.8$) are shown in Figure 7. When the buffer increases, the loss probability decreases very quickly at the beginning, but when it is up to about 3000 bytes, the improvement decreases as the buffer size increases further. In this situation, the number of overdue packets becomes dominant while the number of overflowed packets still decreases further. When the loading is heavy, the above three scheduling schemes cannot serve all the packets before their deadlines, although both E-LOSS and EDF-FB can reduce the overflowed packets. They cannot improve the loss performance further because there is still a proportion of packets that are dropped due to the deadline violation. When the buffer size increases large enough (larger than 6000 bytes as shown in Figure 7), the dropped packets are mainly caused by the delay violation, so all the above three schemes converge to the same loss rate. Similar to EDF-FB, there exists some optimal buffer size at which E-LOSS achieves the highest relative improvement over EDF. This optimal buffer size occurs when E-LOSS can continue to reduce the overflowed packets while the overdue packets are still not dominant.

4.5 Effect of Buffer Size with Heterogeneous Traffic

As mentioned in previous sections, the real-time traffic may have different QoS requirements, so the scheduler is expected to schedule different traffic classes with diverse QoS constraints. In this section, the performances of EDF, EDF-FB and E-LOSS are studied for the heterogeneous traffic input in terms of different delay (d_i) and loss (ϵ_i) requirements. The d_i and ϵ_i of the four traffic classes are selected as shown in Table 1.

The EDF scheduler makes scheduling decisions only based on the deadline stamped in packet headers without considering loss requirements, so the traffic with a more stringent delay constraint has a higher priority to get service and hence obtains a lower loss rate. Meanwhile,

Table 1: QoS parameters of the four traffic classes in the simulation

class i	1	2	3	4
d_i (sec)	0.25	0.3	0.25	0.3
ϵ_i	0.01	0.001	0.001	0.01

by considering the queue length information to reduce the overflowed packets, the EDF-FB can obtain a better performance in terms of loss rate; however, under the heterogeneous traffic condition, loss rate coherence between classes arises so that different classes have similar loss performance. In this section, the performance of the E-LOSS scheduler is studied under the heterogeneous traffic input. Four simulation scenarios with different traffic combinations are studied, in which the total number of sources is kept to be 10 but they belong to different traffic classes and the network topology is kept the same as in Figure 4.

4.5.1 Case 1

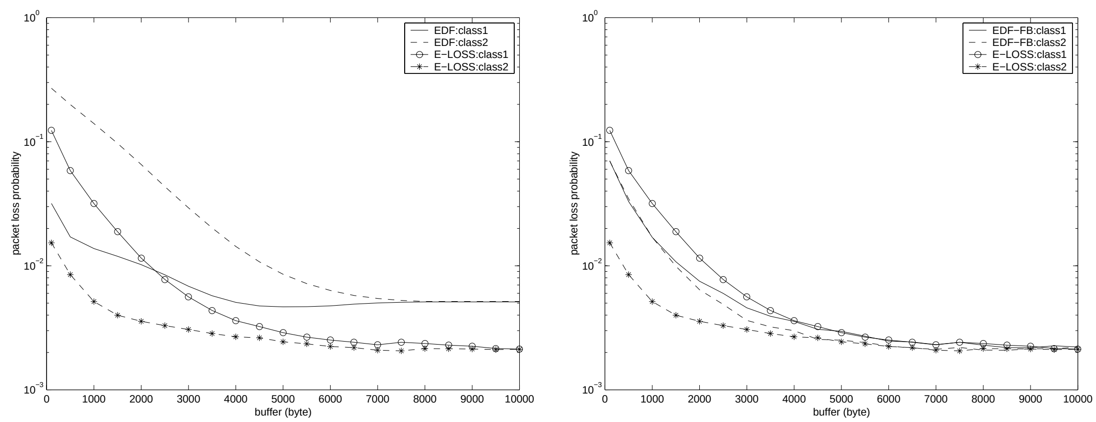

(a) EDF v.s. E-LOSS

(b) EDF-FB v.s. E-LOSS

Figure 8: Loss probability v.s. buffer size under heavy loading and heterogeneous traffic condition of Case 1 ($\rho = 0.8$).

In this case, five sources belong to class 1 while the other five sources belong to class 2. As shown in Figure 8(a), class 1 in EDF has a lower loss rate while its actual loss requirement is rather loose; class 2 in EDF has a higher loss rate while it expects a lower loss rate, so if EDF needs to meet the loss requirements of both classes, the resources for class 1 have to be over-provisioned. As shown in Figure 8(b), both classes have a similar loss rate in EDF-FB. Therefore, class 1 in EDF-FB also needs over-provisioning of resources. On the other hand, the two classes in E-LOSS are differentiated in terms of loss rate. When the buffer increases large enough, both classes have a smaller loss rate than those in EDF.

4.5.2 Case 2

In this case, five sources belong to class 3 while the other five sources belong to class 4. Because the delay requirements for class 3 and class 4 are the same as those of class 1 and class 2 respectively, the loss performance in this case is actually the same as that in Case 1 for both

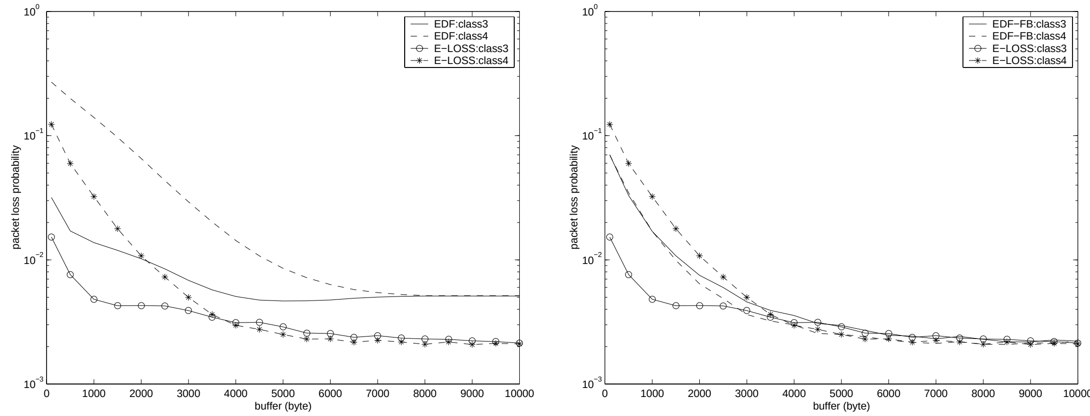

(a) EDF v.s. E-LOSS

(b) EDF-FB v.s. E-LOSS

Figure 9: Loss probability v.s. buffer size under heavy loading and heterogeneous traffic condition of Case 2 ($\rho = 0.8$).

EDF and EDF-FB; however, because the loss requirements for these two classes ($\epsilon_3 = 0.001$ and $\epsilon_4 = 0.01$) are the opposite to those for class 1 and class 2 ($\epsilon_1 = 0.01$ and $\epsilon_2 = 0.001$), so that the loss performance for the two classes in E-LOSS is switched around correspondingly as shown in Figure 9.

4.5.3 Case 3

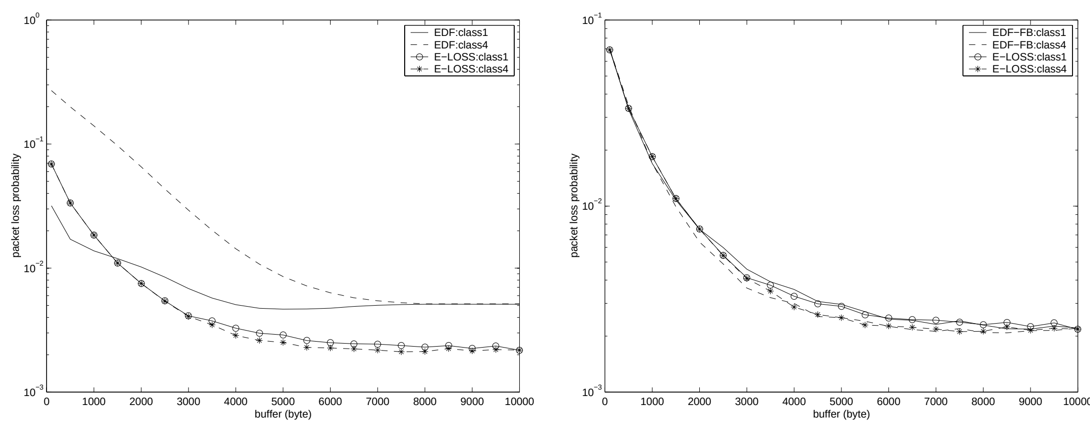

(a) EDF v.s. E-LOSS

(b) EDF-FB v.s. E-LOSS

Figure 10: Loss probability v.s. buffer size under heavy loading and heterogeneous traffic condition of Case 3 ($\rho = 0.8$).

In this case, five sources belong to class 1 while the other five sources belong to class 4. Again, the loss performance in this case is still the same as that in Case 1 in both EDF and EDF-FB. On the other hand, since both classes have the same loose loss requirement, they have a similar loss performance in E-LOSS. In Figure 10(a), the two classes may still have the different loss performances in EDF, even if they have different loss rate requirements. In Figure

10(a), the two classes start to obtain a better performance in E-LOSS than those in EDF when the buffer size is larger than 1500 bytes. Meanwhile, in Figure 10(b), both EDF-FB and E-LOSS have similar performances.

4.5.4 Case 4

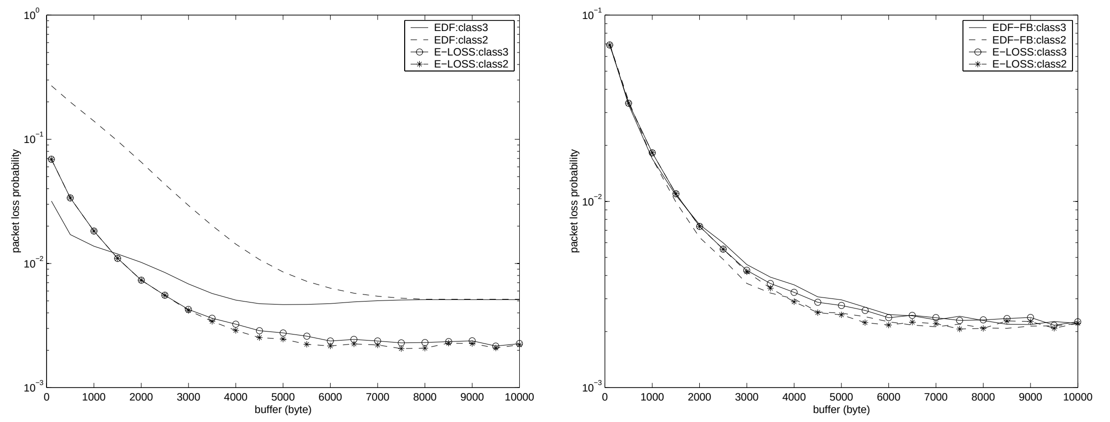

(a) EDF v.s. E-LOSS

(b) EDF-FB v.s. E-LOSS

Figure 11: Loss probability v.s. buffer size under heavy loading and heterogeneous traffic condition of Case 4 ($\rho = 0.8$).

In this case, five sources belong to class 3 while the other five sources belong to class 2. By the same token in case 3, the loss performance in this case is still the same as that in Case 1 in both EDF and EDF-FB. Because class 3 and class 2 have the same loss requirement even though the loss requirement is more stringent than that of class 1 and class 4 in Case 3, E-LOSS in Case 4 shows a similar loss performance as that in Case 3 as shown in Figure 11.

5 Conclusion

With more and more real-time multimedia applications running on the current Internet, the Internet is expected to support packet delivery with a diverse set of QoS guarantees. These applications can usually tolerate some losses without much degradation in quality so that they can be carried as UDP datagrams.

In the hierarchical scheduling framework, the best-effort traffic and real-time traffic share the network resources using the link-sharing scheduler PGPS, while in each class a specific scheduler can be deployed. Due to the fact that different UDP flows may have different QoS requirements so they can be further mapped into sub-classes. Because of the desirable property for EDF that minimizes the maximum lateness of packets, a modified EDF scheme (E-LOSS) is proposed to accomplish packet delivery for the real-time UDP traffic with heterogeneous delay and loss requirements .

Based on the generic EDF scheme together with a Loss Measurement Server (LMS) introduced to perform the on-line loss measurement, E-LOSS aims to maintain a good delay performance while meeting the loss requirement of the traffic. If the queue length is built up, packets belonging to the corresponding flow are likely to be dropped by either deadline violation or buffer overflow. LMS monitors the queues and serves the flow with the maximum normalized loss ratio so that different loss requirements can be supported efficiently. A threshold is set for each queue for LMS to identify whether the queue has been built up. The simulation studies show that in the case of homogeneous traffic, E-LOSS has a similar performance as EDF-FB and outperforms EDF; in the case of heterogeneous traffic, E-LOSS can control the loss rate for the different traffic classes with more flexibility while loss rates in both EDF-FB and EDF are always fixed when delay requirements are readily given.

The major difficulty for E-LOSS is that the current E-LOSS scheme is implemented with the per-flow buffer structure, which may cause the scalability problem when it is deployed in core IP routers. However, the per-class buffer structure may overcome this scalability constraint. Second, the simulation studies assume the ON/OFF traffic model, but recent studies show that the Internet traffic is self-similar [4], [5], so it is necessary to study the scheduler's performance with self-similar traffic models. Besides, the proposed scheduling schemes are only studied in a simple network scenario of one congested link. We also intend to explore these mechanisms in more complicated scenarios with multiple congested links.

References

[1] R. Guerin and V. Peris, "Quality-of-service in packet networks: basic mechanisms and directions", *Computer Networks*, vol.31, no.3, pp.169-189, Feb. 1999.

[2] Rong Pan, Balaji Prabhakar, and Konstantinos Psounis, "CHOKe: a stateless active queue management scheme for approximating fair bandwidth allocation", *Proceedings of INFOCOM'2000*, vol.2, pp.942-951, Mar. 2000.

[3] Jon C. R. Bennett, Hui Zhang, "Hierarchical packet fair queueing algorithms", *IEEE/ACM Transactions on Networking*, vol.5, no.5, pp.675-89, Oct. 1997.

[4] V. Paxson and S. Floyd, "Wide area traffic: the failure of Poisson modeling", *IEEE/ACM Transactions on Networking*, vol.3, no.3, pp.226-44, Jun. 1995.

[5] K. Thompson, G. J. Miller, and R. Wilder, "Wide-area Internet traffic patterns and characteristics", *IEEE Network*, vol.11, no.6, pp.10-23, Nov.-Dec. 1997.

[6] Yinggang Xie and Tao Yang, "Cell Discarding Policies Supporting Multiple Delay and Loss Requirements in ATM networks", *Proceedings of GLOBECOM 97*, vol.2, pp.1075-1080, 1997.

[7] Tao Yang and Jie Pan, "A Measurement-Based Loss Scheduling Scheme", *Proceedings IEEE INFOCOM '96*, vol.3, pp.1062-1071, 1996.

[8] Larry L. Peterson and Bruce S. Davie, Computer networks : a systems approach. 2nd edition, San Francisco : Morgan Kaufmann Publishers, 2000.

[9] A.M.O. Lai and Danny H.K. Tsang, "Service Scheduling Schemes for Finite Buffer in ATM Networks", *International Journal of Distributed System & Networks*, vol.3, no.2, pp.95-107, 2000.

[10] Xiaojun Hei and Danny H.K. Tsang, "The Earliest Deadline First Scheduling with Active Buffer Management for Real-Time Traffic in the Internet", International Conference on Networking (ICN'01), July 9-13, 2001, CREF, Colmar, France.

[11] NS (Network Simulator): http://www.isi.edu/nsnam/ns/.