

ON THE PERFORMANCE OF DIFFERENTIATED SERVICES FOR IP INTERNETWORKS

Chie Dou*, Tian-Shiuh Jeng*
Shu-Wei Wang** and Kuo-Cheng Leu**

**Department of Electrical Engineering
National Yunlin University of Science and Technology
**Computer & Communications Research Laboratories
Industrial Technology Research Institute, Taiwan, R.O.C.*

Abstract Traffic prioritization is an effective, yet relatively simple, tool for providing differentiated services. This paper investigates the performance of different traffic types for IP internetworks with traffic prioritization under various traffic patterns via numerical analysis. The following seven traffic types were considered: Network Control, 'Voice', 'Video', Controlled Load, Excellent Effort, Best Effort and Background, as suggested in IEEE P802.1D. Traffic patterns were appropriately selected to cover a wide range of traffic variations. This paper demonstrated how the performance of individual traffic types is affected by the distribution of server utilization of all types. The results obtained in this paper could also be helpful to network administrators in configuring the "bandwidth reservation" for IP internetworks.

Keywords: Traffic Prioritization, Differentiated Services, IP Internetworks

1. INTRODUCTION

Traditional IP internetworks were originally designed to provide only a "best-effort" delivery service for each application on a first-come, first-serve basis. This has been, and should continue to be, sufficient for the vast majority of network traffic. E-mail, file transfer and even network-based Fax services can be handled in this manner without suffering from its unpredictable (but finite) delay characteristics. However, several classes of application demand a higher level of service. For example [1]:

- New applications such as Voice over IP and Video Webcasting (distance learning, corporate broadcast, etc.) will not tolerate highly variable delays in the delivery of packets, even though the bandwidth they require is modest.

- Existing business-critical applications shifted to the corporate internetwork - for example, mainframe-based on-line transaction processing - may need to be prioritized above other tasks; in many cases, the applications themselves require highly predictable delivery service from the network.

The business and technical issues that network managers will need to understand in order to successfully deliver differentiated service in their converged networks were thoroughly discussed in [2]. Traffic prioritization is an effective, yet relatively simple, tool for providing differentiated services. Traffic from multimedia and business-critical applications should be assigned to higher-priority queues that are given preferential treatment over lower-priority queues and thus receive lower delay and better performance. However, a fundamental problem is arisen with this approach. The problem is that some applications seek to use the entire or a large portion of the available resource; if they use one or more higher-priority queues, they can completely lock out lower-priority queues for a period of time. This can cause lower-priority sessions to be dropped or to be so slow that they are unusable [3]. To ensure that no class achieves more than a given proportion when the line is under stress. Bandwidth can be reserved so that different streams of data are guaranteed a minimum quantity of bandwidth. This feature allows multimedia and business-critical traffic to have low latency while still permitting other network applications to run effectively. Class-based queuing (CBQ) [4] is one of such techniques which deliver differentiated services by reserving network bandwidth on a static (permanent until manually reconfigured) basis. While no class consumes more than a certain bandwidth under any circumstance, another fundamental problem is arisen with this approach. That is the guaranteed percentage of bandwidth for each traffic type needs to be determined appropriately in advance.

To deal with the second problem mentioned above, a prior knowledge about the performance of individual traffic classes under various traffic patterns could be helpful to network administrators in configuring the bandwidth reservation. For example, the network administrators may want to know what is the performance of best-effort and/or business-critical data traffic when the link utilization of multimedia traffic is high. If the multimedia traffic is further divided into voice and video traffic, what is the performance of other data streams when voice or video traffic becomes heavy? For IP internetworks, which adopt traffic prioritization, this paper provides an effective way through numerical analysis to obtain the useful information needed in configuring the bandwidth reservation.

2. TRAFFIC TYPES

As suggested in the Annex of [5], the following list of traffic types were considered: Network Control (NC), 'Voice' (VO), 'Video' (VI), Controlled Load (CL), Excellent Effort (EE), Best Effort (BE), and Background (BK). To fully understand the performance of individual traffic types under various traffic patterns, we assume that each traffic type has its own service queue. Thus, the server utilization of a particular queue can be regarded as the fraction of the bandwidth

occupied by the corresponding traffic type. The bandwidth occupancy of the type is equal to the amount of bandwidth reserved for it if the data of the corresponding traffic type fully utilizes its reserved bandwidth and does not exceed it. To have a prior knowledge about whether a configuration of bandwidth reservation is effective, this paper provides a simple while systematic approach to get the answer. Our approach adopts strict priority queuing algorithm and assumes that all traffic types fully utilize their reserved bandwidths. That is, the server utilization of each priority queue can be used to represent the fraction of bandwidth reserved for the corresponding traffic type. It had already been shown that the average waiting time for each priority queue depends solely on the server utilization of every individual queue in the system. Hence, the numerical results of the waiting time analysis for strict priority queuing systems can be used to evaluate the performance of individual traffic types under various configurations of bandwidth reservation.

3. STRICT PRIORITY QUEUING SYSTEMS

Strict priority queuing algorithm is a simple while effective means to provide useful service differentiation in the IP internetworks. IEEE P802.1D (incorporating IEEE P802.1P) recommended this algorithm as the default algorithm for selecting frames for transmission on each Bridge port. This system is known also by the name of head-of-the-line (HOL) priority queuing or fixed priority queuing, as shown in Fig. 1. In this section we summarize some important results of the queuing analysis derived for the system from [6].




Figure 1 Head-of-the-line priority queue.

We assume that arriving packets belong to one of a set of N different priority classes, indexed by the subscript p ($p = 1, 2, \dots, N$). We adopt the convention that the larger the value of the index associated with the priority group, the higher is the so-called priority associated with that group. We consider a fairly general model based on the system $M/G/1$. Thus we assume that packets from priority group p arrive in a Poisson stream at rate λ_p packets per second; each packet from this group has its service time selected independently from the distribution $B_p(x)$ with mean \bar{x}_p sec. We define the following:

$$\lambda = \sum_{p=1}^N \lambda_p, \quad \bar{x} = \sum_{p=1}^N \frac{\lambda_p}{\lambda} \bar{x}_p, \quad \rho_p = \lambda_p \bar{x}_p, \quad \text{and} \quad \rho = \lambda \bar{x} = \sum_{p=1}^N \rho_p.$$

The interpretation of ρ here is, as usual, the fraction of time the server is busy (so long as $\rho < 1$). Moreover, ρ_p is the fraction of time the server is busy with packets from group p (again for $\rho < 1$).

3.1 Calculating Average Waiting Times

We consider the case of nonpreemptive systems. We study the system from the point of view of a newly arriving packet from priority group p (say); we shall refer to this packet as the “tagged” packet. Let us denote by w_0 the average delay to our tagged packet due to the packet found in service. Since we have a Poisson process, then ρ_i is the probability that our tagged packet finds a type- i packet in service.

With Poisson arrivals, we have

$$W_0 = \sum_{i=1}^N \rho_i \frac{\overline{x_i^2}}{2x_i} \quad (1)$$

where $\overline{x_i^2}$ is the second moment of service time for a packet from group i .

Let W_p be the average waiting time for packets from group p . From [6], we have

$$W_p = W_0 + \sum_{i=p}^N \overline{x_i} \lambda_i W_i + \sum_{i=p+1}^N \overline{x_i} \lambda_i W_p \quad p=1, 2, \dots, N \quad (2)$$

Solving (2) for W_p , we have

$$W_p = \frac{W_0 + \sum_{i=p+1}^N \rho_i W_i}{1 - \sum_{i=p}^N \rho_i} \quad p=1, 2, \dots, N \quad (3)$$

Solving recursively, we obtain the solution

$$W_p = \frac{W_0}{(1 - \sigma_p)(1 - \sigma_{p+1})} \quad p=1, 2, \dots, N \quad (4)$$

where $\sigma_p = \sum_{i=p}^N \rho_i$ and $\sigma_{N+1} = 0$.

From (4), we see the effect of those packets of equal or higher priority present in the queue when our packet arrives as given by the denominator term $1 - \sigma_{p+1}$.

4. NUMERICAL RESULTS

This section evaluates the performance of different traffic types under various traffic patterns in an IP internetwork. Since we assume each traffic type has a corresponding priority queue, the server utilization of each queue can be used to represent the mean offered load of the corresponding traffic type. Thus, a specific traffic pattern can be denoted by a 7-tuple vector $(\rho_1, \rho_2, \rho_3, \rho_4, \rho_5, \rho_6, \rho_7)$, where ρ_i is the server utilization of queue i . Table 1 describes the characteristics of data packets belonging to each individual traffic type. We assume the packet length for each type is uniformly distributed between a minimum packet size (MIN) and a maximum packet size (MAX) of that type. To calculate the first and the second moment of the service time, i.e. \overline{x} and $\overline{x^2}$, for each traffic type, we assume the link capacity $C=1.54$ Mbps. For convenience we further use ρ_1 to denote the aggregate server utilization of the lower three traffic types, including the BK, BE and EE. That

is, $\rho_{II} = \rho_1 + \rho_2 + \rho_3$. Also, we use ρ_{II} to denote the aggregate server utilization of the higher four traffic types, including the CL, VI, VO and NC. That is, $\rho_{II} = \rho_4 + \rho_5 + \rho_6 + \rho_7$. It is evident that larger ρ_I means that the more bandwidth is occupied by the lower priority traffic types. On the contrary, larger ρ_{II} means that the more bandwidth is occupied by the higher priority traffic types. This section investigates the performance of different traffic types by making comparisons between six appropriately selected different traffic patterns.

Table 1

Priority	Traffic type	Packet length (in byte)		\bar{x}	\bar{x}^2
		MIN	MAX		
7	NC	40	60	$2.60 \cdot 10^{-4}$	$5.38 \cdot 10^{-8}$
6	VO	64		$3.32 \cdot 10^{-4}$	$1.10 \cdot 10^{-7}$
5	VI	188		$9.75 \cdot 10^{-4}$	$9.50 \cdot 10^{-7}$
4	CL	150	250	$1.04 \cdot 10^{-3}$	$1.10 \cdot 10^{-6}$
3	EE	200	400	$1.56 \cdot 10^{-3}$	$2.50 \cdot 10^{-6}$
2	BE	400	800	$3.11 \cdot 10^{-3}$	$1.00 \cdot 10^{-5}$
1	BK	500	1500	$5.19 \cdot 10^{-3}$	$2.91 \cdot 10^{-5}$

Note: define

$$a = 8 \cdot \text{MIN}/C, \quad b = 8 \cdot \text{MAX}/C;$$

$$\bar{x} = (a+b)/2; \quad \bar{x}^2 = \int_a^b \frac{1}{b-a} x^2 dx = (a^2 + ab + b^2)/3.$$

Figures 2 - 4 display the numerical results for three different traffic patterns considered in the first group, respectively. The first pattern assumes the server utilization of individual priority queues shares the same percentage of the total server utilization, ρ . That is $\rho_i = \rho/7$ for all i . Fig. 2 shows the average waiting time versus the total server utilization ρ for this case. Two scales are shown for the same case in Fig. 2. The second pattern assumes $\rho_1 = \rho/7$ and $\rho_1 = \rho_2 = \rho_3 = \rho_1/3$; $\rho_{II} = 6\rho/7$ and $\rho_4 = \rho_5 = \rho_6 = \rho_7 = \rho_{II}/4$. The third pattern assumes $\rho_1 = 6\rho/7$ and $\rho_1 = \rho_2 = \rho_3 = \rho_1/3$; $\rho_{II} = \rho/7$ and $\rho_4 = \rho_5 = \rho_6 = \rho_7 = \rho_{II}/4$. Evidently, the second pattern represents the case in which the server utilization is heavily dominated by the data of higher priority traffic types. The third pattern represents the case in which the server utilization is heavily dominated by the data of lower priority traffic types. Figures 3 and 4 show the numerical results of these two cases respectively. The bolded curve shown in each figure represents the corresponding *M/G/I* result of the system without traffic prioritization. Obviously, the *M/G/I* curve shown in Fig. 3 performs better than the corresponding curves shown in Figs. 2 and 4. This is because the packet lengths of higher priority types are smaller than that of lower priority types. In Fig. 3, the server utilization is heavily dominated by the data of higher priority types. But this does not mean that all curves shown in Fig. 3 perform better than the corresponding curves shown in Figs. 2 and 4. From figures

2-4 we observe that the curves of w_1 , w_6 and w_7 shown in Fig. 3 perform better than the corresponding curves shown in Figs. 2 and 4. On the contrary, the curves of w_2 , w_3 , w_4 and w_5 shown in Fig. 3 perform worse than the corresponding curves shown in Figs. 2 and 4 if the value of ρ is larger than some threshold for each type of curves. This observation clearly indicates that if the server is heavily loaded with the data of higher priority types as shown in Fig. 3, the corresponding performance of types including BE, EE, CL and VI becomes worse than the other two cases. But the performance of types BK, VO and NC still performs better than the other two cases. From Fig. 4 we notice that if the server is heavily loaded with the data of lower priority types, the performance of w_4 , w_5 , w_6 and w_7 is very close, and the individual performance of w_2 , w_3 , w_4 and w_5 is the best among three cases.




Figure 2 W_i vs. ρ with traffic pattern $(\frac{\rho}{7}, \frac{\rho}{7}, \frac{\rho}{7}, \frac{\rho}{7}, \frac{\rho}{7}, \frac{\rho}{7}, \frac{\rho}{7})$.




Figure 3 W_i vs. ρ with traffic pattern $(\frac{\rho}{21}, \frac{\rho}{21}, \frac{\rho}{21}, \frac{3\rho}{14}, \frac{3\rho}{14}, \frac{3\rho}{14}, \frac{3\rho}{14})$.




Figure 4 W_i vs. ρ with traffic pattern $(\frac{2\rho}{7}, \frac{2\rho}{7}, \frac{2\rho}{7}, \frac{\rho}{28}, \frac{\rho}{28}, \frac{\rho}{28}, \frac{\rho}{28})$.




Figure 5 W_i vs. ρ with traffic pattern $(\frac{\rho}{70}, \frac{4\rho}{35}, \frac{\rho}{70}, \frac{3\rho}{35}, \frac{3\rho}{35}, \frac{3\rho}{5}, \frac{3\rho}{35})$.




Figure 6 W_i vs. ρ with traffic pattern $(\frac{3\rho}{35}, \frac{24\rho}{35}, \frac{3\rho}{35}, \frac{\rho}{70}, \frac{\rho}{70}, \frac{\rho}{10}, \frac{\rho}{70})$.

Figures 5 and 6 display the numerical results for two different traffic patterns considered in the second group, respectively. The first pattern assumes $\rho_I = \rho/7$ and $\rho_{II} = 6\rho/7$. The second pattern assumes $\rho_I = 6\rho/7$ and $\rho_{II} = \rho/7$. The server utilization of each traffic type for both patterns follows the same distribution: $\rho_1 = \rho/10$, $\rho_2 = 8\rho_1/10$, $\rho_3 = \rho_1/10$, $\rho_4 = \rho_{II}/10$, $\rho_5 = \rho_{II}/10$, $\rho_6 = 7\rho_{II}/10$ and $\rho_7 = \rho_{II}/10$. Similar to the results we have observed from figures 3 and 4, we also notice that the curves of W_2 , W_3 , W_4 and W_5 shown in Fig. 5 perform worse than the corresponding curves shown in Fig. 6 as ρ is getting larger. Since the total server utilization is dominated by the voice (VO) traffic only, we further observe that the curves of W_4 , W_5 and W_6 shown in Fig. 5 move upwards more rapidly than the corresponding curves shown in Fig. 3 as ρ is getting larger. In Fig. 6 we assume the total server utilization is dominated by the best effort traffic. Hence, the curves of W_3 , W_4 , W_5 , W_6 and W_7 are very close, and the performance of W_1 and W_2 is much worse than the corresponding performance shown in Fig. 4. From Fig. 6 we notice that if the total server utilization is dominated by lower priority traffic, the performance of higher priority types is almost not affected by lower priority traffic.

All the three different traffic patterns considered in the third group assume that $\rho_I = \rho/7$ and $\rho_{II} = 6\rho/7$. We further assume all three patterns have $\rho_1 = \rho_I/10$, $\rho_2 = 8\rho_1/10$, and $\rho_3 = \rho_1/10$, the same fraction as used in the previous two patterns. The difference between these three patterns occurred in the distribution of server utilization of higher priority types. For each pattern, ρ_{II} is dominated by a

selected higher priority type with fraction 7/10. The other three higher priority types share ρ_{II} with equal fraction 1/10. The dominant types of three traffic patterns are the controlled load (CL), video (VI) and voice (VO) respectively. Table 2 listed the results of performance comparisons among these three traffic patterns for all priority types. From Table 2 we observe an important property of strict priority queuing systems. This property is that if the total server utilization is dominated by priority type i , then the performance of priority type $i-1$ will be the worst among three patterns.

Table 2

Pattern no.	Traffic type						
	BK(W_1)	BE(W_2)	EE(W_3)	CL(W_4)	VI(W_5)	VO(W_6)	NC(W_7)
1. CL dominant	W	W	⊙	B	B	B	W
2. VI dominant	M	M	M	⊙	M	M	M
3. VO dominant	B	B	B	M	⊙	W	B

(B:Best, M:Middle, W:Worst)

All the three different traffic patterns considered in the fourth group also assume $\rho_1 = \rho/7$ and $\rho_{II} = 6\rho/7$. The difference between these three patterns occurred in the distribution of server utilization of lower three priority types. For each pattern, ρ_1 is dominated by a selected lower priority type with fraction 8/10. The other two lower priority types share ρ_1 with equal fraction 1/10. The dominant types of ρ_1 in three different traffic conditions are the BK, BE and EE respectively. We further assume all three patterns have $\rho_4 = \rho_{II}/10$, $\rho_5 = \rho_{II}/10$, $\rho_6 = 7\rho_{II}/10$ and $\rho_7 = \rho_{II}/10$. Table 3 listed the results of performance comparisons among these three traffic conditions for all priority types. It is evident that for higher four priority types the third pattern (EE dominant) performs the best, and the first pattern (BK dominant) performs the worst. This is because the mean packet sizes of priority types are getting smaller as priority level increases. Smaller packet sizes always result in lower average waiting times. Table 3 clearly shows that the worst cases of BE and BK were EE dominant and BE dominant respectively. This matches the important property of strict priority queuing systems we have described above.

Table 3

Pattern no.	Traffic type						
	BK(W_1)	BE(W_2)	EE(W_3)	CL(W_4)	VI(W_5)	VO(W_6)	NC(W_7)
1. BK dominant	B	B	X	W	W	W	W
2. BE dominant	⊙	M	B	M	M	M	M
3. EE dominant	M	⊙	Y	B	B	B	B

Note: for $\rho > 0.8$, X=M and Y=W;
 $\rho < 0.8$, X=W and Y=M.

Figures 7, 8 and 9 display the numerical results for three different traffic patterns considered in the fifth group, respectively. We assume $\rho_1 = 6\rho/7$ and $\rho_{II} = \rho/7$ for

these three patterns. The distributions of server utilization of all traffic types for these three patterns are the same as the previous three patterns considered in the fourth group. We found that the comparison result for this group of traffic patterns is almost the same as Table 3 with some minor exceptions in BK and EE types. Although the respective curves belonging to two different groups are unlike in appearance.




Figure 7 W_i vs. ρ with traffic pattern $(\frac{24\rho}{35}, \frac{3\rho}{35}, \frac{3\rho}{35}, \frac{\rho}{70}, \frac{\rho}{70}, \frac{\rho}{10}, \frac{\rho}{70})$.




Figure 8 W_i vs. ρ with traffic pattern $(\frac{3\rho}{35}, \frac{24\rho}{35}, \frac{3\rho}{35}, \frac{\rho}{70}, \frac{\rho}{70}, \frac{\rho}{10}, \frac{\rho}{70})$.




Figure 9 W_i vs. ρ with traffic pattern $(\frac{3\rho}{35}, \frac{3\rho}{35}, \frac{24\rho}{35}, \frac{\rho}{70}, \frac{\rho}{70}, \frac{\rho}{10}, \frac{\rho}{70})$.

All the three different traffic patterns considered in the sixth group also assume $\rho_I = 6\rho/7$ and $\rho_{II} = \rho/7$. The distributions of server utilization of all traffic types for these three patterns are the same as those considered in the third group. The distribution of server utilization for lower three priority types follows $\rho_1 = \rho_I/10$, $\rho_2 = 8\rho_I/10$ and $\rho_3 = \rho_I/10$. For each pattern, ρ_{II} is dominated by a selected higher priority type with fraction $7/10$. The other three higher priority types share ρ_{II} with equal fraction $1/10$. Since the packet sizes of higher priority types are small

and their aggregate utilization ρ_{II} only occupies a small portion of the total server utilization, the corresponding performance of W_1 , W_2 and W_3 in all three traffic conditions is almost the same. Different distributions of server utilization for higher priority types affect the performance of W_4 , W_5 , W_6 and W_7 very slightly.

5. CONCLUSIONS

Numerical analysis for the strict priority queuing systems was used to investigate the performance of individual traffic types under various traffic patterns in IP internetworks. Traffic patterns were appropriately selected to cover a wide range of traffic variations, and were classified into six groups. Performance comparisons and discussions were made for each group of patterns. Some remarkable observations are listed as follows:

- (1) If the server is heavily loaded with the data of traffic type i , then the performance of traffic type $i-1$ will become worse dramatically. This will also result in the degradation of the performance of traffic types with priority lower than $i-1$ to some degree.
- (2) If the server is heavily loaded with the data of voice (VO) type, the performance of all traffic types will become worse obviously except NC.
- (3) Since the mean packet sizes of traffic types are getting larger as priority level decreases, the performance of higher priority types is affected by the distribution of server utilization of lower priority types. Traffic patterns with EE dominant perform better than those with BE and/or BK dominant.
- (4) If the server is heavily loaded with the data of lower priority types, the performance of higher priority types can be affected by the traffic of lower priority types only slightly.
- (5) Since the packet sizes of higher priority types are rather small, if the aggregate utilization ρ_{II} only occupies a small portion of the total server utilization, the performance of individual traffic types will be affected by the distribution of server utilization of higher priority types very slightly.

References

- [1] --"Implementing Guaranteed Service in IP Internetworks," *A technical paper from Torrent Networking Technologies*. <http://www.torrentnet.com>
- [2] Chuck Semeria and Frank Fuller, "3Com's Strategy for Delivering Differentiated Service Levels," 3Com Systems Marketing, Feb. 6, 1998.
- [3] *Cisco White Paper*, "Interface Queue Management," Posted: Dec. 5, 1997. http://www.cisco.com/warp/public/614/quemg_wp.htm
- [4] Ashley Stephenson, "Class-based Queuing: Managing Broadband Access to the Internet," *Network World*, Vol. 14, Issue 21, May 26, 1997.
- [5] ISO/IEC Final DIS 15802-3, Information Technology - local and metropolitan area networks - Common specifications -- Part 3: Media Access Control (MAC) Bridges, May 25, 1998.
- [6] Leonard Kleinrock, *Queueing Systems, Volume II: Computer Applications*, John Wiley & Sons, Inc., 1977.