Computer Communications 24 (2001) 1661–1672
computer communications
www.elsevier.com/locate/comcom

Research into two source-based control algorithms for Internet traffic flows

Chyan Yang*, Chen-Hua Fu

Institute of Information Management, National Chiao Tung University, MB 307, 1001 Ta Hsueh Road, Hsin Chu 300, Taiwan, ROC

Received 17 April 2000; revised 9 January 2001; accepted 11 January 2001

Abstract

TCP and UDP are the major applications over the Internet, the characteristics of them lead to different network transmission behaviors. Two source-based mechanisms are proposed in this paper to regulate TCP and UDP flows. One is the congestion control mechanism, which uses TCP flows' congestion signal to regulate the flows at the source node. The other is the time slot mechanism, which is a time-sharing application to control their flow transmission. Based on the priorities of flows, different bandwidth proportions are allocated and differential services are provided for flows. Several scenarios are simulated to observe the transmission operations of these two mechanisms. Simulation results show some insights into two mechanisms. Moreover, the several simulation parameters that may impact the performance of these two mechanisms are summarized.

Keywords: Differential service; Network congestion; Congestion control mechanism; Time slot mechanism; Priority; Transmission performance

1. Introduction

The TCP and UDP are the two major protocols over the Internet. These two protocols have different transmission operations. TCP is connection orientated whereas UDP is connectionless. TCP uses 'slow start' mechanism [2] as the end-to-end congestion control mechanisms to prevent network congestion. UDP simply uses 'store and forward' mechanism to transfer its data. If there is no proper control mechanism to handle UDP's transmissions, UDP traffic will share most of the bandwidth over the Internet. These characteristics of TCP and UDP lead to different network transmission behaviors. The drop-and-run UDP is unfavorable to self-controlled TCP when the proportion of UDP is relatively higher than that of TCP.

Since most of the Internet applications are based on TCP, the performance of TCP will impact the Internet efficiency. How to improve the TCP transmission performance and restrict too much bandwidth shared by UDP transmissions is the focus of this study. In this study, a differential service mechanism is proposed to handle the transmission of TCP and UDP traffic. Priority setting for each TCP or UDP transmission ensures each flow gets the different bandwidth share. The higher the priority is, the better the performance gets.

Two flow control mechanisms are proposed to control the TCP and UDP traffic transmissions in this study. One is a congestion control mechanism. It uses the TCP traffic congestion signal and the priority of each flow to control the transmissions. The other is a time slot control mechanism. According to the priority of each flow, the different amount of time slots is assigned. Network simulator ns-2 [1] is used to simulate the traffic transmission operations of these two mechanisms. Simulation results indicate that both of the two control mechanisms can provide a differential service for the larger transmission data, such as 1 MB, and the congestion control mechanism has a better transmission performance than the time slot mechanism. Moreover, some factors, such as queueing disciplines and parameter settings, may impact the performance of the proposed control mechanisms.

2. Analysis of TCP/UDP traffic transmission

The transmission protocols of TCP and UDP are quite different. The difference between TCP and UDP calls for different transmission behaviors and, therefore, they have different transmission performance over the Internet.

TCP is a connection-oriented protocol and it uses the 'slow start' mechanism [2] to control the traffic transmission

* Corresponding author. Tel.: +886-3-5712121, ext. 57407/57417; fax: +886-3-5723792.

E-mail addresses: cyang@cc.nctu.edu.tw (C. Yang), u8634804@cc.nctu.edu.tw (C.-H. Fu).

0140-3664/01/$ - see front matter
PII: S0140-3664(01)00311-5

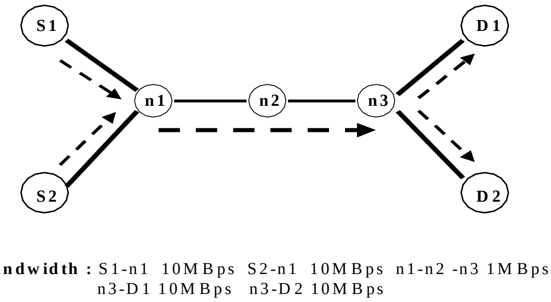

Fig. 1. A topology of simulation scenario.

and response the network congestion. UDP is a connectionless protocol and it has no congestion control mechanism to response to network congestion. When the network is congested, an extreme unfairness situation exists among TCP/UDP flows: the bandwidth should be shared by the TCP flows, which are occupied by the UDP flows. This unfairness results from responsive and irresponsible flows competing for bandwidth. The UDP flow effectively 'shuts out' the responsive TCP traffic [3].

In this study, two simple simulation scenarios are established with the ns-2 network simulator [1] to observe the transmission situation of TCP and UDP traffic. The network topology is shown as Fig. 1. The S1, S2, D1 and D2 are the source and destination nodes. There are four TCP flows and three UDP flows from S1 to D1, and three TCP flows and UDP flows from S2 to D2. Each TCP flow uses FTP traffic as its traffic source, whereas the traffic source of each UDP flow is the constant bit rate traffic.

The first simulation scenario is that each TCP/UDP flow transmits the assigned data size simultaneously. The length of assigned data size is from 10 KB to 100 MB. The simulation results of the transmission time of each TCP and UDP flow are listed in Table 1. Table 1 shows that all transmission time of the UDP flows are shorter than the TCP flows and the transmission performance differences between TCP and UDP flows increase as the transmission size increases. This phenomena shows that the UDP flows always have larger shares of bandwidth than that of TCP flows and therefore receive a better transmission performance. On the other hand, TCP flows regulate their transmission by the 'slow start' mechanism when the network is congested; they do not transmit packets until UDP flows finish their transmission. This is a reason why the TCP flows have the longer transmission time than the corresponding UDP flows.

The second simulation scenario is that the TCP/UDP flows have a same period to transmit traffic and record the traffic sizes they transmitted. This simulation tries to record the transmission size of the TCP/UDP flows in a same period. The range of a simulation transmission period is 1–10000 s and nine timer checkpoints are selected to collect the transmission size of the TCP/UDP flows. Fig. 2 shows the results of this simulation; the ratio variability of bandwidth shared by TCP flows and UDP flows.

Table 1
A listing of transmission performance of TCP/UDP flows (Unit: second)

Flows		Size				
		10 KB	100 KB	1 MB	10 MB	100 MB
S1 to D1	TCP1	2.6	13.5	122.9	1080.1	10462.1
	TCP2	2.7	14.7	121.7	1066.7	10475.8
	TCP3	2.7	16.4	207.5	1089.9	7488.9
	TCP4	2.7	14.8	121.8	1073.1	10436.5
	UDP1	0.8	5.0	48.5	474.7	4769.6
	UDP2	1.4	6.0	48.0	480.8	4794.8
	UDP3	0.7	5.3	48.1	481.0	4802.3
S2 to D2	TCP5	2.6	14.5	122.8	1079.1	10482.2
	TCP6	2.7	14.6	124.8	1080.7	9851.1
	TCP7	2.8	15.5	129.9	710.3	7544.9
	UDP4	0.8	4.2	46.5	479.2	4797.0
	UDP5	0.6	3.2	47.8	480.5	4782.4
	UDP6	0.9	5.1	49.2	471.6	4788.6

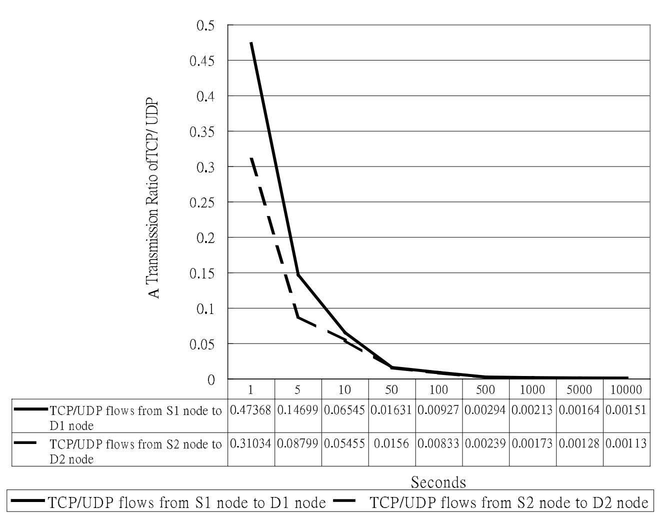

Fig. 2. Variability of TCP/UDP transmission ratio.

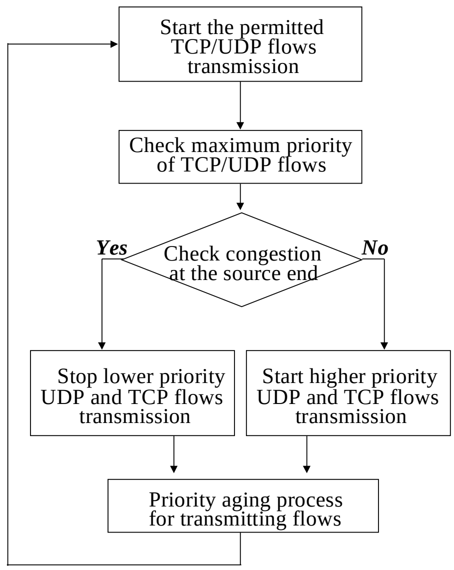

Fig. 3. A diagram of congestion control mechanism.

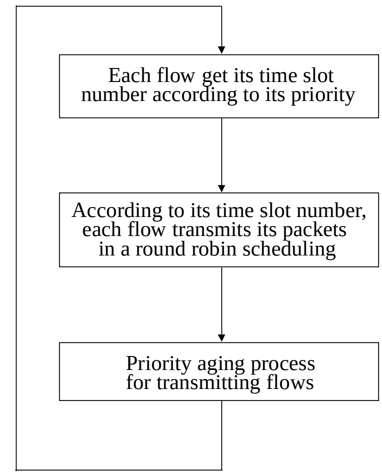

Fig. 4. A diagram of time slot mechanism.

The curves in Fig. 2 show the ratio of bandwidth shared by TCP flows and UDP flows that is inversely changed as the transmission period increases. The curves also show that the bandwidth share of TCP flows drop dramatically from the first second to the 50th second, and the bandwidth share of TCP flows is less than 1% after the 100th second. These results show exactly why the UDP flows have the better performance than the TCP flows in the first simulation scenario, because most of bandwidth is shared by the UDP flows during the transmission period.

The results from the second simulation demonstrate that the TCP flows have a higher bandwidth share at the beginning of a transmission period because the network is not so congested. When the UDP flows continue to send their packets on the network, gradually the network congestion becomes even more congested. The TCP flows continue to responsd the network congestion by reducing their transmission rates, and their bandwidth shares continue to decrease. If the UDP flows keep their traffic transmission, UDP traffic will receive more bandwidth share. This in turn will cause TCP flows to continue slowing down TCP transmissions and the bandwidth share of TCP is going lower than that of UDP.

The above two simulation results show that an unfairness situation exists between the transmissions of the TCP flow and the UDP flow. So it is obvious that a proper control mechanism is needed to regulate the transmission of UDP flow and prevent the bandwidth share overused by the UDP flows. This is why the two source-based [4] flow control mechanisms are studied here.

3. Source-based flow control mechanisms

Generally speaking, traffic management or bandwidth control is a mechanism that treats different users or flows differently — ranging from simple Weighted Fair Queueing to RSVP and per-session traffic scheduling.[1,2] Most proposed differential service mechanisms involve the capability of the gateways on a routing path and these mechanisms need gateway devices to support their operations.

In this study, two source-based flow control mechanisms are compared: the congestion control mechanism and the time slot mechanism. Both mechanisms use a priority of flows to regulate the utilization of network bandwidth. Priority overwrites the types of protocols. That is, the priority 1 of UDP traffic takes higher preference to the priority 2 of TCP traffic. For the same type of protocol, priority determines the preference.

3.1. Congestion control mechanism

The congestion control mechanism is a source-based flow control mechanism. It operates at the source node, when it detects the network congestion, it reacts with a slower transmission, hoping not to worsen the network congestion. The idea of this control mechanism comes from the characteristics of TCP and UDP flows and the simulation results from [3]. If the congestion signal from TCP flows can be used as a congestion indicator for the source node, this could help the source node control the TCP and UDP traffic transmissions. When the transmission path is congested, the source node can stop the transmissions of lower priority flows and let higher priority flows keep their transmissions. With regulated transmission, higher priority flows can have better transmission performance.

Depending on the importance and time constraint of a transmission, network administrators may assign a proper transmission priority to TCP/UDP flow at the source nodes. A time critical flow can receive a higher transmission priority than otherwise. The congestion control mechanism

[1] URLhttp://www.ietf.org/html.charters/diffserv-charter.html.
[2] URLhttp://diffserv.lcs.mit.edu/.

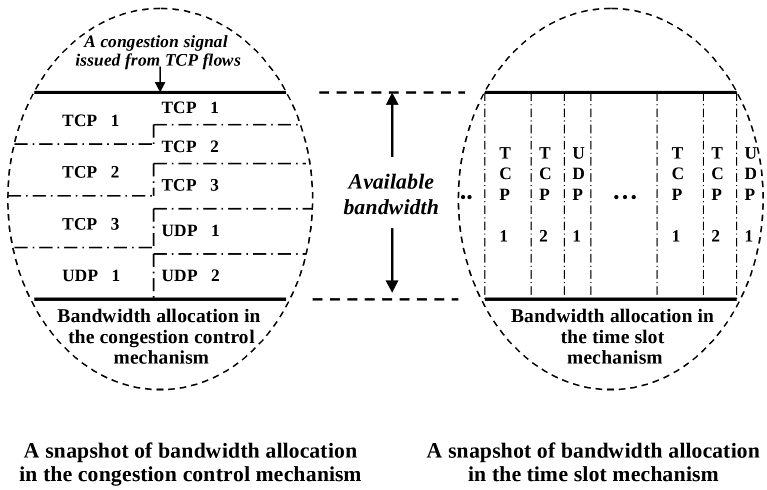

Fig. 5. Bandwidth allocations of congestion control and time slot mechanisms.

collects flows' priority information. TCP and UDP flows can concurrently transmit their packets. The diagram of the congestion control mechanism is shown in Fig. 3.

3.2. Time slot mechanism

The time slot mechanism is a time-sharing application that divides bandwidth into many transmission units (time slot). Each transmission unit is a time slot. In each time slot, the source node only allows one flow to transmit its packets and this flow can use all available bandwidth as much as it can. All other flows must yield the right of way. How many time slots can a flow get? It depends on the priority, assigned by network administrators at a source node, of a flow.

A round robin scheduler is used by the time slot mechanism to arrange each flow's transmission. A high priority flow receives more time slots than otherwise and therefore receives the required bandwidth. In other words, the UDP flows can no longer occupy the bandwidth share irresponsibly. Moreover, the transmission performance of each flow can be ensured with its priority. The diagram of time slot mechanism is shown in Fig. 4.

Both of the congestion control and time slot mechanisms can regulate the TCP and UDP flow transmissions. Fig. 5 is a snapshot of bandwidth utilization with these two mechanisms.

4. A simulation of congestion control and time slot mechanisms

Several simulation scenarios are simulated to illustrate the operations of these two mechanisms. With the simulation results, one can obtain some transmission performance about these two mechanisms. We also investigated the factors that may affect the algorithms.

The Network Simulator — ns (version 2 beta release 5) [1,5] is used as the simulation tool. The simulation topology is same as Fig. 1. A number of TCP/UDP flows are simulated to transmit packets from the S1 and S2 source nodes to the D1 and D2 destination nodes. Fig. 6 shows the transmission procedure in each simulation scenario.

The ratio of Internet TCP/UDP flow is basic to our simulation scenarios. From the MCI/NSF's very-high-performance Backbone Network Service (vBNS) project,[3] one can find the ratio of TCP/UDP flows is 90:10. Based on this, with 100 flows, the TCP may vary from 81 to 99 whereas UDP varies from 19 to 1 during simulation. With different TCP/UDP flow combinations, the transmission performance of the proposed mechanisms can be analyzed further. The transmission size is another factor that may impact the transmission behavior of flows. The 10 KB and 1 MB are the two transmission lengths that are used as the sources of TCP/UDP flows in the simulation scenarios.

For a differential service simulation, six bits in IP header are reserved to indicate the priority of the flow [6,7]. These six levels of priorities are available in both proposed mechanisms. For the congestion control mechanism, priority is used to determine whether a flow continues its transmission when the network is congested. For the time slot mechanism, priority is used to determine the number of time slots allocated for a flow. Let P denote the priority of a flow where $P = 1, 2, \ldots, 6$. Let $\mathrm{tsn}(P)$ denote the number of time slots assigned to a flow with priority P. A binary bandwidth allocation of $\mathrm{tsn}(P)$ can be defined as $\mathrm{tsn}(P) = 2\hat{}(9 - P)$. Therefore, the difference between $\mathrm{tsn}(i)$ and $\mathrm{tsn}(j)$ is $|2^i - 2^j|$. The time slot mechanism assigns 256 time slots as the round robin transmission time to the highest priority flow ($P = 1$). A lowest priority flow ($P = 6$) will receive only 8 time slots as its round robin transmission time. In this research, four level priorities are assigned to the TCP flows. Let TCP_i denote a TCP flow of priority i. Two priorities, 3 and 4, are assigned to the UDP flows. Let UDP_i denote a UDP flow of priority i.

Queueing disciplines are also important since they may

[3] URLhttp://www.vbns.net/stats/flows/html/index.html.

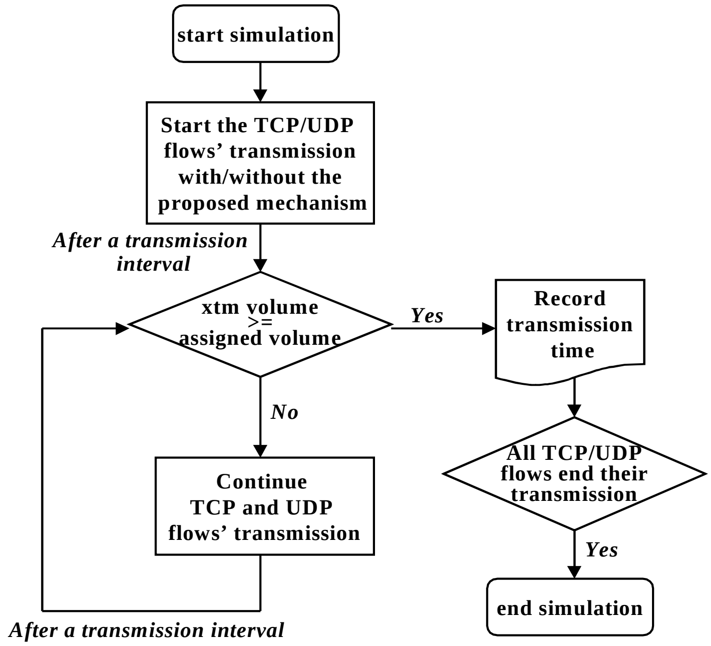

Fig. 6. A flow chart of simulation scenario.

impact the transmission performance of the two proposed control mechanisms. In our simulations, four different queueing disciplines: FCFS (First Come First Serve), SFQ (Stochastic Fair Queue) [8], RED (Random Early Detection) [9] and DRR (Deficit Round Robin) [10] are implemented to schedule network applications' transmissions.

The congestion control mechanism and time slot mechanism may coexist with other transmission control mechanisms and with the best-effort traffic. Co-working with a different control mechanism, the proposed mechanisms may have a different transmission performance. Table 2 illustrates the five different simulation environments. Table 3 summarizes key parameters in the simulation scenarios.

5. Results and analysis

Transmission size, queueing disciplines, TCP/UDP ratios, environment settings and parameter settings are the major parameters used for the simulation. Based on these parameters, several simulation results are observed and analyzed.

5.1. Sensitivity to transmission size

Two different sizes of data transmission, 10 KB and 1 MB, are used to measure the size sensitivity of the congestion control mechanism and the time slot mechanism. The simulation results show the congestion control mechanism provides a differential service for 10 KB TCP/UDP flows. But, there are cases that the transmission performance is inconsistent with their transmission priorities. For 1 MB TCP/UDP flows, the congestion control mechanism provides a significant differential service among the TCP/UDP flows. In cc/cc and cc/be transmission environments, most of the 1 MB TCP/UDP flows receive a different transmission performance based on their transmission priorities. That is, the priority of the flow indicates the transmission performance.

Likewise, the differential service transmission behaviors only occur in the 1 MB TCP/UDP flows in the ts/ts and ts/be transmission environments. But, when the transmission size

Table 2
Specification of the five transmission environment settings

Let cc(S_i, D_j) denote the congestion control mechanism applied to the flows from source S_i to destination D_j; Let ts(S_i, D_j) denote the time slot mechanism applied to the flows from source S_i to destination D_j; Let be(S_i, D_j) denote best-effort traffic applied to the flows from source S_i to destination D_j

Environment settings	Purpose
1. cc(S_1, D_1) and cc(S_2, D_2): cc/cc	To simulate a transmission environment controlled by the congestion control mechanism
2. ts(S_1, D_1) and ts(S_2, D_2): ts/ts	To simulate a transmission environment controlled by the time slot mechanism
3. cc(S_1, D_1) and be(S_2, D_2): cc/be	To simulate a transmission environment controlled by a congestion control mechanism and a best-effort traffic mechanism
4. ts(S_1, D_1) and be(S_2, D_2): ts/be	To simulate a transmission environment controlled by a time slot mechanism and a best-effort traffic mechanism
5. be(S_1, D_1) and be(S_2, D_2): be/be	To simulate a transmission environment controlled by the best-effort traffic mechanism

Table 3
A summary of key simulation parameters

Parameter name	Number of alternatives	
1. Transmission size	2 (10 KB, 1 MB)	2×5×4 = 40
2. Environment setting	5 (see Table 2)	
3. Queueing discipline	4 (FCFS, SFQ, RED, DRR)	
4. TCP/UDP ratio	19 (81:19,…,99:1)	19×2×100 = 3800
5. Traffic sources	2	
6. TCP/UDP flows	100 (81 + 19, …,99 + 1)	

is of 10 KB, a first-come-first-serve transmission behavior exists in each TCP/UDP flow's transmission in the ts/ts and ts/be environments. This first-come-first-serve transmission behavior of the time slot mechanism causes the first transmitted flow to receive the best transmission performance since there is no waiting. However, the last transmitted flow gets the worst transmission performance. Fig. 7 shows that the total transmission time required is linearly proportional to the number of flows. In the ts/ts setting, since the flows start to transmit at the equal spacing of time slots, various ratios of TCP/UDP produce same straight lines and overlays as in Fig. 7.

Fig. 7 also shows that the ts/ts and ts/be settings cannot work properly to provide a differential service for TCP/UDP flows when the transmission size is relatively small to each assigned time slot. Comparing with the congestion control mechanism, the time slot mechanism is more sensitive to a transmission size.

5.2. *Differential service operations in the control mechanisms*

Table 4 shows the summary of average transmission performance of flows from the S1 source node to the D1 destination node and the S2 source node to the D2 destination node. These flows are cc, ts or be. In each of the environment settings, four queueing disciplines (FCFS, DRR, RED and SFQ) are used. The transmission size of each flow is of 1 MB. Fig. 8 shows the average transmission performance of TCP/UDP flows with the FCFS queueing discipline in different transmission environments. Table 4 and Fig. 8 show differential service operations that are provided by the congestion control and time slot mechanisms and the higher transmission priority flows get better transmission performance than the lower transmission priority flows.

Fig. 8 also shows the simulated results of various ratios of TCP/UDP data flows. The line charts in Fig. 8 demonstrate the variations of transmission performance of TCP/UDP flows as the TCP flow number increases and the UDP number decreases. From the transmission performance line charts' fluctuations, there is no obvious evidence showing there is a relationship between a ratio of TCP/UDP flows and their transmission performance. A ratio of TCP/UDP does not impact a differential service operation that is supported by the congestion control and time slot mechanisms.

5.3. *Transmission performance of the control mechanisms*

Table 4 clearly shows that the congestion control mechanism in each case outperforms the time slot mechanism. This simulation result may result from the following three reasons: (1) in the ts/ts environment, the time slot mechanism might suffer from underutilization since the required burst bandwidth is smaller than the time slot assigned; (2) most of the bandwidth in the ts/be environment

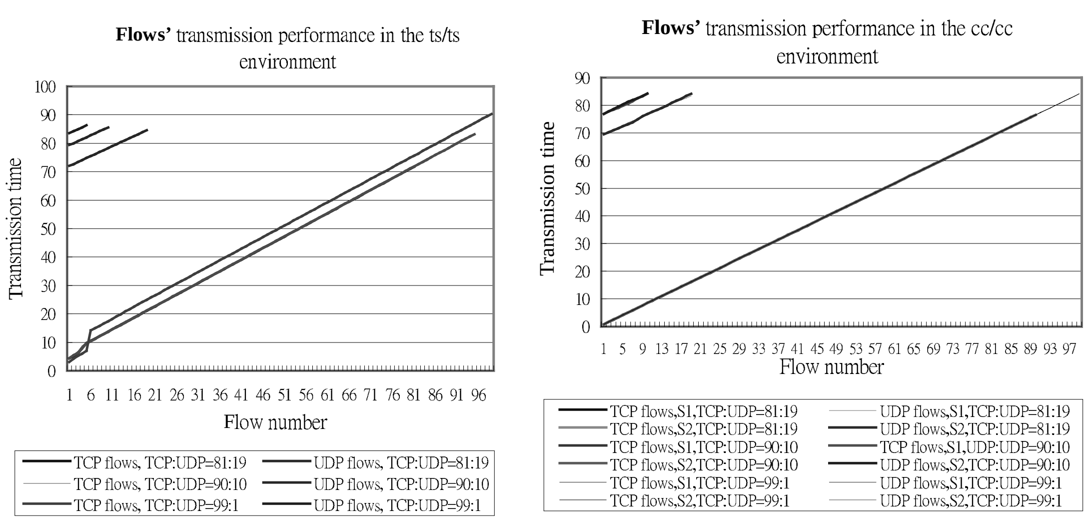

Fig. 7. A performance of time slot mechanism in 10 KB file size.

Table 4
Average TCP/UDP flows' transmission performance
1. In cc/cc, ts/ts, cc/be and ts/be, flows in source traffic 1 are regulated by control algorithms; 2. In be/be, flows in source traffic 1 are best effort flows; 3. In cc/cc and ts/ts, flows in source traffic 2 are regulated by control algorithms; 4. In cc/be, ts/be and be/be, flows in source traffic 2 are best effort flows; 5. Unit: seconds

Transmission environment	Queueing discipline	Flows' priority											
		Source traffic 1						Source traffic 2					
		TCP_1	TCP_2	TCP_3	TCP_4	UDP_3	UDP_4	TCP_1	TCP_2	TCP_3	TCP_4	UDP_3	UDP_4
cc/cc (a pure congestion control mechanism environment)	FCFS	290.39	612.04	1084.67	1581.06	802.65	1390.16	303.58	688.02	1228.05	1661.27	961.70	1450.27
	DRR	250.98	735.68	1626.65	2459.73	1197.29	2037.99	264.20	790.87	1623.42	2423.65	1191.06	2033.07
	RED	258.35	653.51	1244.00	1678.98	1025.00	1484.29	260.94	596.36	1068.91	1541.42	856.28	1326.45
	SFQ	276.19	588.35	1201.76	1697.74	1050.77	1602.53	285.77	680.83	1236.64	1724.79	1103.02	1613.02
ts/ts (a pure time slot mechanism environment)	FCFS	450.42	1114.75	1447.57	1910.88	1533.65	1719.07	450.72	1114.60	1452.96	1910.29	1515.99	1723.58
	DRR	451.61	1115.23	1438.84	1889.38	1527.89	1726.63	451.81	1118.65	1442.59	1893.43	1535.82	1731.57
	RED	454.44	1147.63	1479.98	1928.71	1421.80	1653.31	459.33	1140.56	1481.73	1927.36	1425.11	1653.11
	SFQ	451.09	1114.12	1438.51	1902.18	1529.15	1748.05	451.29	1117.36	1449.24	1880.19	1556.94	1756.44
cc/be (a mixed congestion control mechanism environment)	FCFS	582.19	1232.33	1585.42	1937.91	1402.76	1758.36	754.87	722.30	767.47	752.26	80.17	84.51
	DRR	579.35	1979.88	2679.95	3318.67	2305.03	3047.99	993.82	1003.11	1003.19	1079.70	88.41	89.87
	RED	579.52	1158.97	1551.12	1932.87	1387.72	1748.04	625.81	634.59	642.28	641.01	81.02	84.95
	SFQ	483.67	1187.39	1613.88	1994.12	1486.93	1941.99	718.41	709.83	722.19	704.19	148.70	148.03
ts/be (a mixed time slot mechanism environment)	FCFS	1214.91	1787.47	2265.74	2840.17	1741.27	2058.57	593.68	588.24	586.53	590.25	79.95	83.20
	DRR	1048.57	1637.40	2145.01	2687.73	1743.89	2078.98	761.24	774.40	743.56	800.43	90.62	92.28
	RED	1178.34	1750.06	2237.68	2760.54	1704.01	2015.05	528.78	533.03	526.32	532.72	80.41	84.41
	SFQ	1180.96	1768.99	2272.46	2880.33	1745.17	2045.72	559.36	560.88	556.00	537.61	152.49	151.82
be/be (a pure best-effort traffic environment)	FCFS	1004.34	1030.48	1000.94	1029.12	145.64	151.82	1111.01	1118.17	1126.82	1141.71	139.29	146.16
	DRR	1907.32	1997.61	1924.30	2023.77	157.10	173.75	1626.96	1707.24	1614.93	1706.87	164.29	164.99
	RED	930.06	947.05	890.14	975.90	159.14	166.73	961.78	957.40	984.74	996.37	158.44	166.88
	SFQ	1068.02	1060.22	1056.63	1049.85	240.06	277.87	1086.83	1056.28	1057.76	1052.27	257.96	256.47

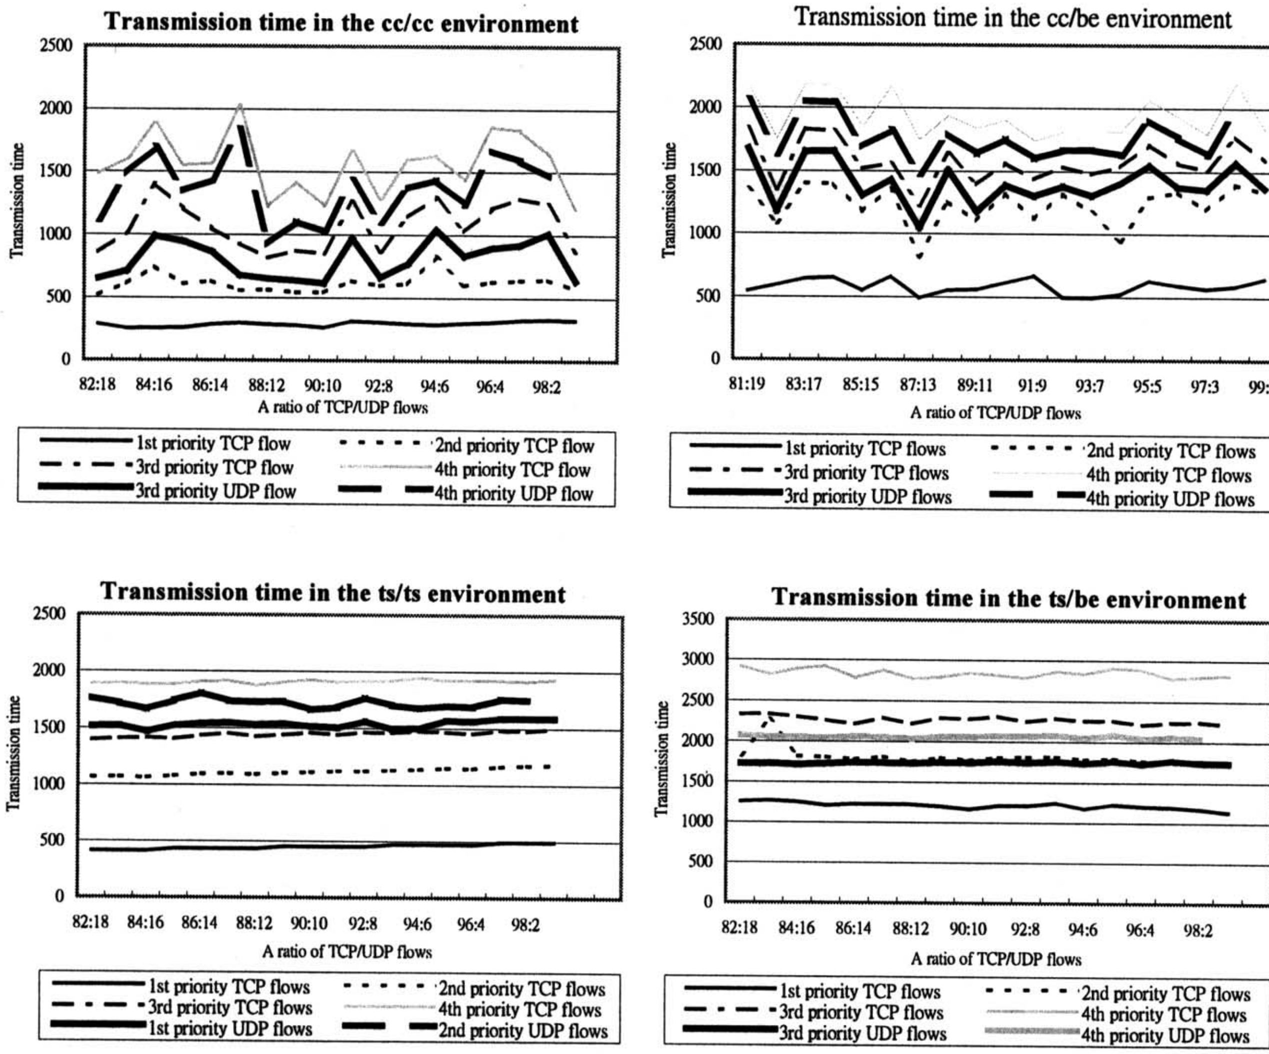

Fig. 8. Differential service operations of the cc/cc, cc/be, ts/ts and ts/be environments.

may be taken by best-effort traffic, only a little bandwidth available for the time slot mechanism to regulate the flows' transmission; (3) the congestion control mechanism always keeps higher priority flows to take the bandwidth whenever possible.

Fig. 9 shows flows' transmission performance in the cc/cc, ts/ts, cc/be and ts/be environments. In Fig. 9, the thin line charts represent the transmission performance of cc/cc and cc/be flows and the thick line charts represent the transmission performance of ts/ts and ts/be flows. From these two group line charts, they clearly show the fluctuations of transmission performance of flows with the different TCP/UDP ratios. Although there is no significant relationship between transmission performance and a TCP/UDP ratio, Fig. 9 shows the transmission performance fluctuations of the cc/cc and cc/be flows are larger than ts/ts and ts/be flows. Therefore, this means that the time slot mechanism is relatively more stable and predictable than the congestion control mechanism with the different TCP/UDP ratios, although the time slot mechanism is slower than the congestion control mechanism.

In addition, Table 4 also shows that UDP_3 has better performance than that of TCP_3, and UDP_4 has better performance than that of TCP_4. With the same transmission priority, UDP flows would outperform TCP flows. Therefore, by assigning a lower priority to UDP flows, when they are not time critical, assures that TCP flows transmit before UDP flows.

5.4. Transmission priority level of the flows

The transmission performance of these two control mechanisms in the cc/be and ts/be environments is interesting. From the cc/be and ts/be row blocks of columns in the Source traffic 1 in Table 4, one can find that the congestion control mechanism has a better transmission performance

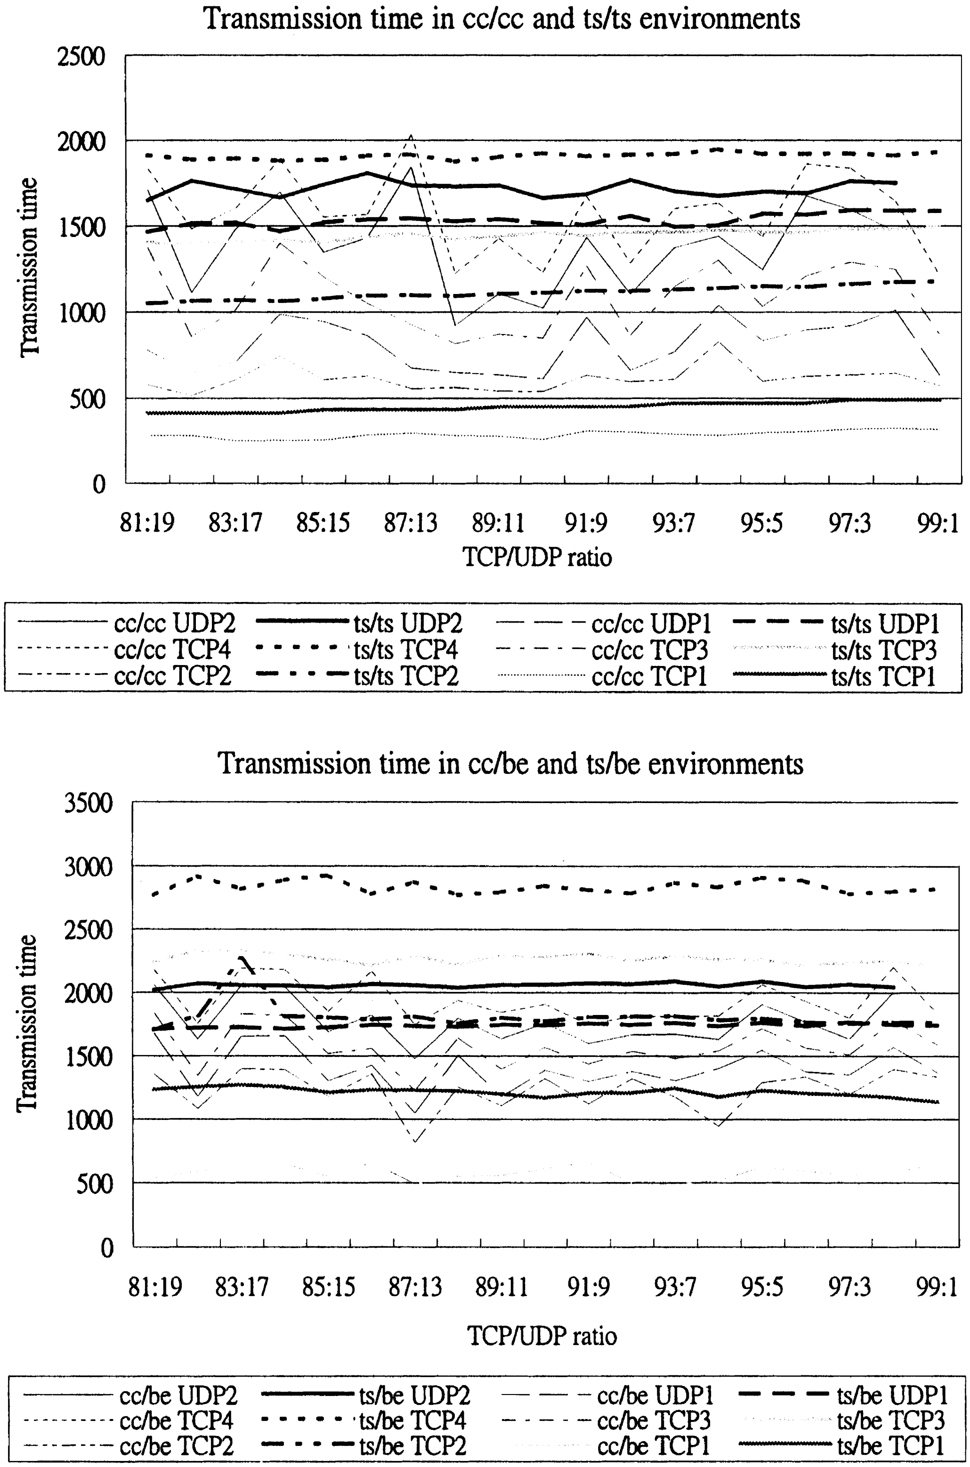

Fig. 9. Flows' transmission performance in different transmission environments.

than that of the time slot mechanism when they operate with best-effort flows. Except that the cc/be TCP_1 flows are better than all the be/be TCP flows, the transmission performance of cc/be TCP_2, TCP_3 and TCP_4 is worse than the corresponding be/be TCP flows. Only the cc/be TCP_1 flows get a guaranteed service. In other words, there is no need to have too many priority levels in a differential service mechanism. Two levels are enough. A high priority flow will receive a guaranteed service and has a better performance than flows of lower priorities. On the other hand, the low priority or the best-effort flows will compete for bandwidth left by flows of the highest priority.

Table 5
Observations between transmission performance and queueing disciplines

Queueing disciplines	Environments	Flows	Performance
FCFS	cc/be, ts/be, be/be	All be UDP flows	Best
DRR	cc/cc, be/be	All cc TCP_2, TCP_3, TCP_4, UDP_3, UDP_4 flows	Worst
		All be TCP_2, TCP_3, TCP_4, UDP_3, UDP_4 flows	
	ts/be	All ts TCP flows	Best
		All be TCP flows	Worst
	cc/cc, ts/ts, cc/be, ts/be	All cc and ts flows	ts flows are better than cc flows
RED	cc/be, ts/be, be/be	All TCP flows	Better
	cc/be, ts/be, be/be	All be TCP flows	Best
	cc/cc, ts/ts, cc/be, ts/be	All UDP flows	Better
	ts/ts	All TCP flows	Worst
		All UDP flows	Best
SFQ	cc/cc, ts/ts	All UDP flows	Worst
	cc/be, ts/be, be/be	All UDP flows	Worst

5.5. A relationship among the control mechanisms, transmission performance and queueing disciplines

Table 4 shows that the transmission performance is not too sensitive to the queueing disciplines. The queueing disciplines, however, do impact the transmission performance when different priorities are imposed on the traffic. Several observations from Table 4 are listed in Table 5.

Table 5 shows there is no significant relationship among control mechanisms, transmission performance and the queueing disciplines. But several simulation results are interesting: (1) all cc/be, ts/be, be/be UDP flows with the FCFS queueing discipline receive the best transmission performance; (2) the DRR queueing discipline favors ts TCP flows and does not favor cc and be TCP flows. The cc and be TCP flows, except TCP_1 flows, get the worst transmission performance in the cc/cc and be environments; (3) the RED queueing discipline favors all cc/be, ts/be and be/be TCP flows. The UDP flows with the RED queueing discipline also receive good transmission performance; (4) all UDP flows with the SFQ queueing discipline have poor transmission performance.

Table 6 shows the standard deviation of transmission time corresponding to flows shown in Table 4. By examining the column source traffic 1 in Table 6, one can find the following observations: (1) the standard deviation value of transmission performance of the cc/cc and cc/be environments are greater than that of the ts/ts and ts/be environments in most cases. The transmission performance of time slot mechanism is more stable and predictable than the transmission performance of congestion control mechanism; (2) in the cc/cc, ts/ts, cc/be and ts/be environments, one also finds that the standard deviation value of the lower priority flows is larger than that of the higher priority flows. That is, the higher priority flows have a more stable and predictable transmission performance than the lower priority flows; (3) the FCFS and RED queueing disciplines might support a more stable transmission performance than other queueing disciplines in the cc/cc environment; (4) with the DRR queueing discipline, all cc/cc, cc/be and be/be TCP flows have an unstable transmission performance; (5) the SFQ queueing discipline could support the most stable transmission performance for all cc/be, ts/be and be/be UDP flows. Moreover, the cc/be and ts/be TCP flows with SFQ queueing discipline have more stable transmission performance than other queueing disciplines.

5.6. Parameter settings of the control mechanisms

Proper parameter settings are important to the control mechanisms simulations; it allows the mechanisms have better control. Three key parameters are investigated: round robin transmission time, number of time slots assigned to each priority and priority aging.

The round robin transmission time setting is important in the time slot mechanism. A flow's round robin transmission time depends on the number of time slot. A proper setting of time slot number benefits both TCP and UDP flows. If the number of time slot is too large, it will lead to a first-come-first-served transmission behavior just as the case shown in Fig. 7. If the number of time slot is too small, a short round robin transmission time is shorter than the transmission round trip time. If a TCP flow cannot get its destination ACKs, the retransmission of a TCP flow will occur repeatedly and deteriorate its transmission performance.

Additionally, if differences among round robin transmission time of different priority flows are too large, a low priority flow may starve. Otherwise, the differential services are not significant among different priority flows. For the time slot mechanism, a time slot number assignment to each

Table 6

Standard deviation of TCP/UDP flows' transmission performance

1. In cc/cc, ts/ts, cc/be and ts/be, flows in source traffic 1 are regulated by control algorithms; 2. In be/be, flows in source traffic 1 are best effort flows; 3. In cc/cc and ts/ts, flows in source traffic 2 are regulated by control algorithms; 4. In cc/be, ts/be and be/be, flows in source traffic 2 are best effort flows; 5. Unit: seconds

Transmission environment	Queueing discipline	Flows' priority											
		Traffic source 1						Traffic source 2					
		TCP_1	TCP_2	TCP_3	TCP_4	UDP_3	UDP_4	TCP_1	TCP_2	TCP_3	TCP_4	UDP_3	UDP_4
cc/cc (a pure congestion control mechanism environment)	FCFS	22.517	72.429	199.938	241.349	148.040	260.736	19.667	89.446	196.809	218.273	202.538	248.081
	DRR	17.258	112.100	313.526	381.021	263.624	402.590	16.700	202.259	331.519	370.441	285.425	372.951
	RED	16.538	86.521	239.346	262.605	263.354	270.553	13.773	123.985	226.898	235.124	207.280	254.656
	SFQ	15.343	93.327	307.992	316.727	314.763	328.746	16.436	176.114	229.475	273.345	248.735	280.365
ts/ts (a pure time slot mechanism environment)	FCFS	27.919	39.522	30.074	18.342	37.844	43.350	27.919	39.191	29.401	16.292	54.336	39.405
	DRR	27.999	39.377	29.616	28.104	54.286	41.867	27.999	39.877	29.087	20.891	50.322	34.866
	RED	25.765	38.634	34.566	26.690	57.413	36.154	33.895	37.015	26.838	29.518	59.306	38.962
	SFQ	27.955	39.456	29.285	38.297	53.982	38.232	27.955	39.889	30.439	29.491	26.245	12.653
cc/be (a mixed congestion control mechanism environment)	FCFS	57.343	161.123	184.227	170.965	171.031	156.047	104.166	79.531	119.173	100.266	43.787	41.791
	DRR	82.326	189.004	254.046	341.464	227.877	263.819	137.282	132.209	146.395	140.546	27.010	25.289
	RED	79.373	181.252	178.283	177.972	189.809	184.889	64.801	56.536	55.394	69.448	43.992	41.861
	SFQ	36.247	180.497	152.119	153.947	147.039	121.845	58.455	49.443	45.035	41.996	14.786	14.630
ts/be (a mixed time slot mechanism environment)	FCFS	33.812	28.092	38.032	51.455	16.024	17.573	36.756	33.581	29.353	31.377	43.644	40.908
	DRR	61.972	53.980	34.040	50.787	101.308	62.246	135.728	136.586	139.961	119.835	26.628	24.589
	RED	45.988	21.583	36.292	45.825	18.054	20.184	25.924	17.051	27.381	32.826	43.979	41.851
	SFQ	21.151	20.548	16.173	33.261	42.455	23.705	36.456	31.500	31.187	34.650	14.002	13.858
be/be (a pure best-effort traffic environment)	FCFS	91.992	95.689	88.126	109.32	85.349	81.708	89.560	98.907	82.092	90.487	81.233	77.753
	DRR	286.383	250.773	262.411	262.335	62.128	49.027	229.266	279.655	232.57	258.491	57.219	55.615
	RED	72.812	90.609	75.115	73.894	87.951	83.680	75.985	78.553	72.776	64.789	87.444	83.869
	SFQ	86.711	61.566	87.851	86.688	37.404	33.572	110.706	103.583	98.643	93.170	41.795	43.027

priority is important. The binary bandwidth allocation guarantees that high priority flows receive better performance than low priority flows. Meanwhile, this allocation scheme does not starve the low priority flows. The optimal bandwidth allocation of round robin transmission time deserves a further study.

With numerous simulations, one can find that a setting of priority aging time is important for both control mechanisms. A short priority aging time allows the lower priority flows to be upgraded sooner than otherwise. In that case, soon all the flows become the highest priority. This traffic pattern in turn degenerates into a best-effort traffic and the differential service is not supported any more. A long priority aging time, however, may cause a lower priority flow to starve because other higher priority flows may keep coming and jump ahead of the queue.

6. Conclusion

TCP flows are responsive to the network congestion whereas UDP flows are irresponsible. UDP flows often deteriorate the network congestion, sometimes even cause a congestion collapse. The uncontrollable UDP flow is a major problem over the Internet. The congestion control mechanism and time slot mechanism are the two source-based flow control mechanisms studied in this research. These two mechanisms are applied at the source node to regulate the transmissions of TCP and UDP flows. With these two control mechanisms, UDP flows are regulated and are not irresponsible to the network congestion. Moreover, based on the priorities of flows, differential services can be provided by these two mechanisms. High priority flows would receive a better transmission performance than low priority flows.

These two source-based control mechanisms regulate the TCP and UDP flows at the source node. They are compatible with the current transmission operation environment over the Internet. No additional device, protocol, or control mechanism is needed to implement these two mechanisms. The only operation cost of these two mechanisms is the execution time at the source node.

Acknowledgements

This research is sponsored in part by NSC 89-2416-H-009-011

References

[1] S. McCanne, Sally Floyd, ns-Network Simulator, URLhttp://www-mash.cs.berkeley.edu/ns/.
[2] Van Jacobson, Michael J. Karels, Congestion avoidance and control, Proc.ACM SIGCOMM (1988).
[3] Sally Floyd, Kevin Fall, Promoting the use of end-to-end congestion control in the Internet, IEEE/ACM Trans. Network. (1999).
[4] David D. Clark, Wenjia Fang, Explicit allocation of best effort packet delivery service, ACM Trans. Network. 6 (4) (1998) 362–373.
[5] Kevin Fall, Kannan Varadhan, ns Notes and Documentation, May 11, 1999.
[6] Andrew S. Tanenbaum, Computer Networks, 3rd ed., Prentice-Hall, New Jersey, 1996.
[7] Gary R. Wright, W. Richard Stevens, TCP/IP Illustrated, The Implementation, vol. 4, Addison-Wesley, Reading, MA, 1995
[8] Paul E. McKenney, Stochastic fairness queueing, Internetworking: Research and Experience, vol. 2, 1991, pp. 113–131.
[9] Sally Floyd, Van Jacobson, Random early detection gateway for congestion avoidance, IEEE/ACM Trans. Network. 1 (4) (1993) 397–413 URL http://www-nrg.ee.lbl.gov/nrg-papers.html.
[10] M. Shreedhar, George Varghese Lambda, Efficient Fair Queueing using Deficit Round Robin, November 18, 1994.