INTERNATIONAL JOURNAL OF COMMUNICATION SYSTEMS
Int. J. Commun. Syst. 2001; **15**:13–22 (DOI: 10.1002/dac.520)

A comparison of congestion control and time slot algorithms in Internet transmission performance

Chyan Yang[1,*,†] and Chen-Hua Fu[2]

[1]*National Chiao Tung University, Institute of Information Management, Mb307, 1001 Ta Hsueh Road, Hsin Chu 300, Taiwan*
[2]*National Defense Management College, Graduate School of National Defense Information,*
P.O. Box 90046-15, Chung-Ho, Taipei 235, Taiwan

SUMMARY

The characteristics of TCP and UDP lead to different network transmission behaviours. TCP is responsive to network congestion whereas UDP is not. This paper proposes two mechanisms that operate at the source node to regulate TCP and UDP flows and provide a differential service for them. One is the congestion-control mechanism, which uses congestion signal detected by TCP flows to regulate the flows at the source node. Another is the time-slot mechanism, which assigns different number of time slots to flows to control their flow transmission. Based on the priority of each flow, different bandwidth proportions are allocated for each flow and differential services are provided. Simulation results show some insights of these two mechanisms. Moreover, we summarize the factors that may impact the performance of these two mechanisms. Copyright © 2001 John Wiley & Sons, Ltd.

KEY WORDS: differential service; network congestion; congestion-control mechanism; time-slot mechanism; flow priority

1. INTRODUCTION

TCP and UDP are the two major protocols over the Internet. These two protocols have different traffic transmission operations. TCP is connection orientated whereas UDP is connectionless. These characteristics of TCP and UDP lead to different network-transmission behaviours. Since most of the Internet applications are based on TCP, the performance of TCP will impact on the Internet efficiency. The focus of this study is how to improve the TCP transmission performance and restrict the excessive bandwidth taken by UDP transmissions.

Note that there are about a dozen internet-drafts and RFCs related to our subject using the term of 'differentiated services', 'quality of service' or various types of 'forwarding' behaviours [1,2]. In the near future, however, DiffServ-aware devices will still be rare [3]. This is why this

---

* Correspondence to: Chyan Yang, National Chiao Tung University, Institute of Information Management, Mb307, 1001 Ta Hsueh Road, Hsin Chu 300, Taiwan.
† E-mail: cyang@cc.nctu.edu.tw

Received December 1999
Revised October 2000
Accepted February 2001

research focuses on source-based traffic flow control mechanisms to regulate the coexistence of differential service and regular traffic flows.

In this study, two source-based traffic flow control mechanisms are proposed: one is the congestion-control mechanism and the other is the time-slot mechanism. These two mechanisms operate at the source node to regulate TCP and UDP traffic flows. They allocate different bandwidth proportions to different traffic flows according to their priorities. Priority overwrites the types of protocols. That is, the priority 1 of UDP traffic takes higher preference over the priority 2 of TCP traffic. For the same type of protocol, priority determines the preference. With these two control mechanisms, the transmissions of TCP and UDP flows can be regulated and the differential services can be provided at the source node to enhance the transmission performance of higher priority traffic flows.

1.1. Congestion control mechanism

The congestion control mechanism is a source-based traffic flow-control mechanism. 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. The congestion-control mechanism collects the priority information of flows. TCP and UDP flows can concurrently transmit their packets. This mechanism will routinely check the congestion signal issued by TCP flows to detect congestion on the transmission path. If the network is congested, it stops the transmissions of lower priority flows to release some bandwidth share for higher-priority flows to get a better transmission performance. Otherwise, if there is no congestion, perhaps due to more bandwidth available on the network, it starts the transmission of higher-priority flows to enhance the bandwidth utilization. To prevent the transmission starvation of lower-priority flows, this mechanism uses a priority ageing method to upgrade the priority of lower-priority flows. After a transmission period elapses, lower-priority flows can have the higher priority and allocate more bandwidth share to transmit packets.

1.2. Time-slot mechanism

The time-slot mechanism is an application of the time-sharing concept. The bandwidth is divided into many transmission units. 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 or as much as it can. All other flows must yield the right of way during the time slot. The number of time slots that a flow can get depends on the priority, assigned by network administrators at a source node, to a flow.

This strongly ensures that a high-priority flow receives the required bandwidth. With the time-slot control mechanism, the transmission behaviours of TCP and UDP flows will be regulated. UDP flows can no longer occupy the bandwidth share irresponsibly. Moreover, the transmission performance of each flow can be ensured with its priority. A round-robin scheduler is used by the time-slot mechanism to arrange each flow's transmission. The time-slot mechanism

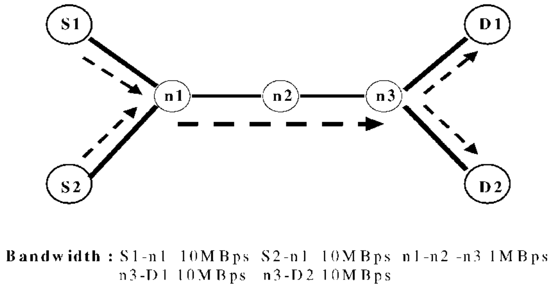

Figure 1. A topology of simulation scenario.

adopts the first-come-first-served principle to append a flow to a round-robin scheduling queue and transmit its packets by turns. When a flow takes turns at transmitting its traffic, the time-slot mechanism assigns a round-robin transmission time to the flow according to its priority. Then a transmission token is assigned to the flow to start its transmission. With a time-slot mechanism, although each flow can get an assigned period to send its packets, a transmission-starvation situation may happen to the lowest-priority flow when higher-priority flows continue to arrive. A priority-ageing method is also incorporated to such a delayed flow.

2. A SIMULATION OF CONGESTION-CONTROL AND TIME-SLOT MECHANISMS

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

The topology of the simulation is shown in Figure 1. TCP/UDP traffic flows are simulated to transmit packets from the S1 and S2 nodes, all traffic flows have the same routing path and share the same bandwidth from N1 node to N4 node, then reach the D1 and D2 nodes. The ratio of Internet TCP/UDP traffic flow is basic to our simulation scenarios. From the MCI/NSF's very high performance Backbone Network Service (vBNS) project [4], one can find that the ratio of TCP and UDP traffic flows is 90 : 10. Based on this, with 100 traffic flows the TCP may vary from 81 to 99 whereas UDP varies from 19 to 1 during simulation. The transmission size is another factor that may impact the behaviour of traffic flows. We use a 10 kbyte file as the smaller traffic-flow source and a 1 Mbyte file as the larger traffic-flow source on the network.

For assigning priority, six bits in the TCP header are reserved to indicate the priority of the flow [5,6]. A bit in a different position represents a different-level priority. The leftmost bit is the highest priority and the rightmost bit is the lowest priority. There are six levels of priorities 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. If the network congests, the flows of lower priority yield way to the flows of higher priority. 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 $P = 1, 2, \dots, 6$. When $P = 1$, the flow is the highest priority, whereas when $P = 6$ the flow is of the lowest priority. Let $\mathrm{tsn}(P)$ denote the number of time slots assigned to a flow with priority P. A binary bandwidth allocation of tsn(P) can be defined as

$$\mathrm{tsn}(P) = 2\hat{}(9 - P)$$

Therefore, the difference between tsn(i) and 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 eight time slots as its round-robin transmission time. In this research, four priority levels 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 impact on the transmission performance of the two proposed control mechanisms. In our simulations, four different queueing disciplines: first come first serve (FCFS), stochastic fair queue (SFQ) [7], random early detection (RED) [8] and deficit round robin (DRR) [9] are implemented to schedule network applications' transmissions.

Various control mechanisms can be used in transmission over the Internet. 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 and behaviours. To simulate different combinations, two groups of end-to-end traffic transmissions with the individual flow-control mechanism are investigated. 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, 'be' represents typical traffic of sources not employing fairness techniques. The five different transmission environments are illustrated in Table I.

3. RESULTS AND ANALYSIS

Parameters used for the simulation include transmission size, queueing disciplines, transmission performance, TCP/UDP ratios, environment setting and parameter settings.

Table I. A specification of the five transmission environment settings.

Environment settings	Purpose
1. cc(S_1, D_1) and cc(S_2, D_2): cc/cc	Congestion-control mechanism
2. ts(S_1, D_1) and ts(S_2, D_2): ts/ts	Time-slot mechanism
3. cc(S_1, D_1) and be(S_2, D_2): cc/be	A congestion-control mechanism and a best-effort traffic mechanism
4. ts(S_1, D_1) and be(S_2, D_2): ts/be	A time-slot mechanism and a best-effort traffic mechanism
5. be(S_1, D_1) and be(S_2, D_2): be/be	Best-effort traffic mechanism

3.1. Sensitivity to transmission size

The simulation results show that the congestion-control mechanism provides a significant differential service among the TCP/UDP flows at the transmission size of 1 Mbyte. Most of the 1 Mbyte TCP/UDP flows receive different transmission performance based on their transmission priorities. In cc/cc and cc/be transmission environments, when the transmission size is large, the priority of the flow dictates the transmission performance. For 10 kbyte TCP/UDP flows, the congestion-control mechanism also provides a differential service. But, there are cases where the transmission performance is inconsistent with their transmission priorities. Some traffic flows without higher priority showed better transmission performance.

Likewise, the differential service transmission behaviours only occur in the 1 Mbytes TCP/UDP flows for ts/ts and ts/be settings. When the transmission size is 10 kbytes, however, the performance of each TCP/UDP flow behaves as a first-come-first-served transmission for both settings. Except for the first flow, to start transmission immediately, all subsequent data flows must wait for their turns even though there are available fragments within under-utilized time slots.

3.2. Differential service operations in the control mechanisms

Table II shows the summary of average-transmission performance of traffic flows from the S1 source node to the D1 destination node. These flows are cc, ts or be. In each environment settings, four queueing disciplines (FCFS, DRR, RED and SFQ) are used. The transmission size of each flow is 1 Mbyte.

Examining Table II, different queueing disciplines do not show significant difference in performance. For both ts and cc mechanisms, performance of a data flow is only dictated by its transmission priority. Figure 2 shows the average transmission performance of TCP/UDP flows with the FCFS queueing discipline in the different transmission environments.

3.3. A relationship between the control mechanisms and queueing disciplines

Table II 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 occur from Table II:

(1) The RED queueing discipline has a better transmission performance for the TCP flows in the cc/cc and cc/be environment. Moreover, the RED queueing discipline does not favour the TCP flows in the ts/ts environment. On the contrary, RED favours the UDP flows.
(2) With DRR-queueing discipline, the lower-priority TCP/UDP flows get worst transmission performance in the cc/cc and cc/be environments.
(3) In the ts/ts and ts/be environments, the four different queueing disciplines do not show much impact on the TCP/UDP flows' transmission performance.
(4) In the different transmission environments with the FCFS- and SFQ-queueing disciplines, the fluctuation of average transmission performance for TCP/UDP flows is smaller than that of the RED and DRR queueing disciplines.

3.4. Transmission performance of the control mechanisms

Table II clearly shows that the congestion-control mechanism in each case outperforms the time-slot mechanism. The time-slot mechanism might suffer from underuse since the required

Table II. Average TCP/UDP flows' transmission performance.

		Flows' priority					
Transmission environment	Queueing discipline	TCP_1	TCP_2	TCP_3	TCP_4	UDP_3	UDP_4
cc/cc (a pure congestion-control mechanism environment)	FCFS	290.3938	612.0431	1084.668	1581.066	802.6508	1390.159
	DRR	250.9803	735.6805	1626.651	2459.728	1197.285	2037.992
	RED	258.3529	653.5088	1244.003	1678.981	1024.996	1484.285
	SFQ	276.186	588.3483	1201.762	1697.743	1050.774	1602.533
ts/ts (a pure time-slot mechanism environment)	FCFS	450.4158	1114.748	1447.569	1910.884	1533.645	1719.071
	DRR	451.6105	1115.231	1438.838	1889.376	1527.888	1726.631
	RED	454.4448	1147.633	1479.975	1928.714	1421.804	1653.313
	SFQ	451.0875	1114.123	1438.505	1902.181	1529.149	1748.045
cc/be (a mixed congestion-control mechanism environment)	FCFS	582.1937	1232.331	1585.42	1937.909	1402.76	1758.358
	DRR	579.3471	1979.884	2679.948	3318.665	2305.033	3047.991
	RED	579.5187	1158.971	1551.121	1932.87	1387.724	1748.035
	SFQ	483.6662	1187.389	1613.882	1994.119	1486.928	1941.988
ts/be (a mixed time-slot mechanism environment)	FCFS	1214.912	1787.466	2265.737	2840.167	1741.265	2058.565
	DRR	1048.573	1637.398	2145.013	2687.725	1743.887	2078.976
	RED	1178.339	1750.063	2237.683	2760.541	1704.006	2015.052
	SFQ	1180.956	1768.989	2272.462	2880.333	1745.173	2045.721
be/be (a pure best-effort traffic environment)	FCFS	1004.341	1030.484	1000.936	1029.117	145.6406	151.82
	DRR	1907.321	1997.608	1924.303	2023.772	157.1049	173.7465
	RED	930.0582	947.0496	890.1357	975.9018	159.142	166.7296
	SFQ	1068.017	1060.224	1056.629	1049.845	240.0575	277.8726
							Unit (s)

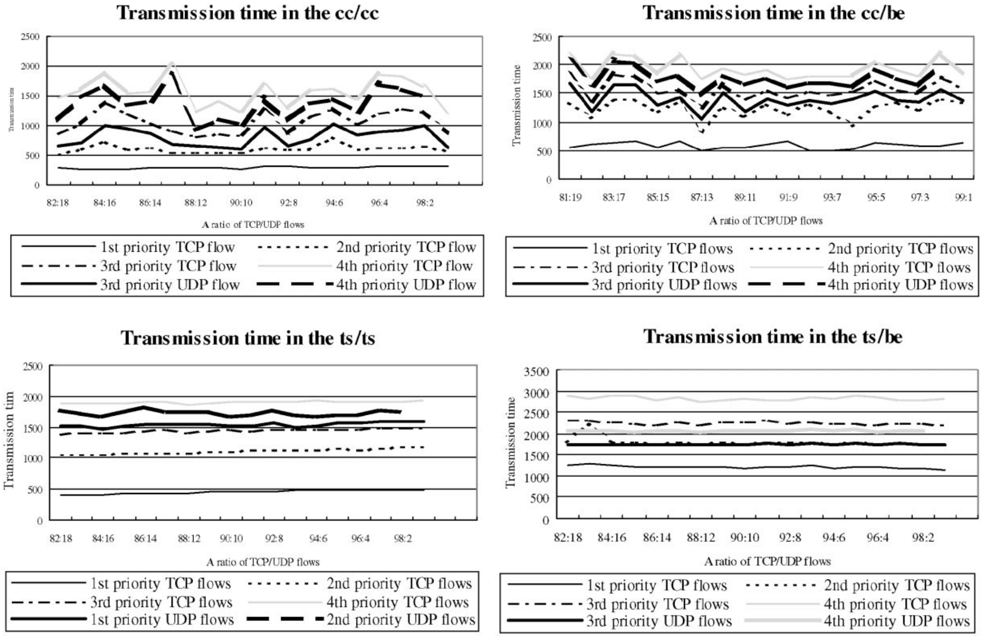

Figure 2. Differential service operations of the cc/cc, cc/be, ts/ts and ts/be environment.

burst bandwidth is smaller than the time slot assigned in the ts/ts environment. Moreover, most of the bandwidth in the ts/be environment may be taken by best-effort traffic, leaving only a little bandwidth available for the time-slot mechanism to regulate the flows' transmission. On the other hand, the congestion-control mechanism always keeps higher priority traffic flows to take the bandwidth whenever possible. This may be another reason why the congestion-control mechanism outperforms the time-slot mechanism in 1 Mbyte simulations.

Table II also shows that UDP_3 has better performance than that of TCP_3, and UDP_4 has better performance than that of TCP_4. Therefore, by assigning a lower priority to UDP flows when they are not time critical, ensures that TCP flows transmit before UDP flows.

The transmission performance of these two control mechanisms in the cc/be and ts/be environments is interesting. From the 3rd and 4th row blocks in Table II, one can find that the congestion-control mechanism has a better transmission performance than that of the time-slot mechanism when they operate with best-effort traffic flows. In contrast to the cc/be TCP_1 flows being better than all the be/be TCP flows, the transmission performance of TCP_2, TCP_3 and TCP_4 in the cc/be is worse than the corresponding TCP flows in the be/be environment. 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 priority. On the other hand, the low priority or the best-effort traffic flows will compete for bandwidth left by flows of the highest priority.

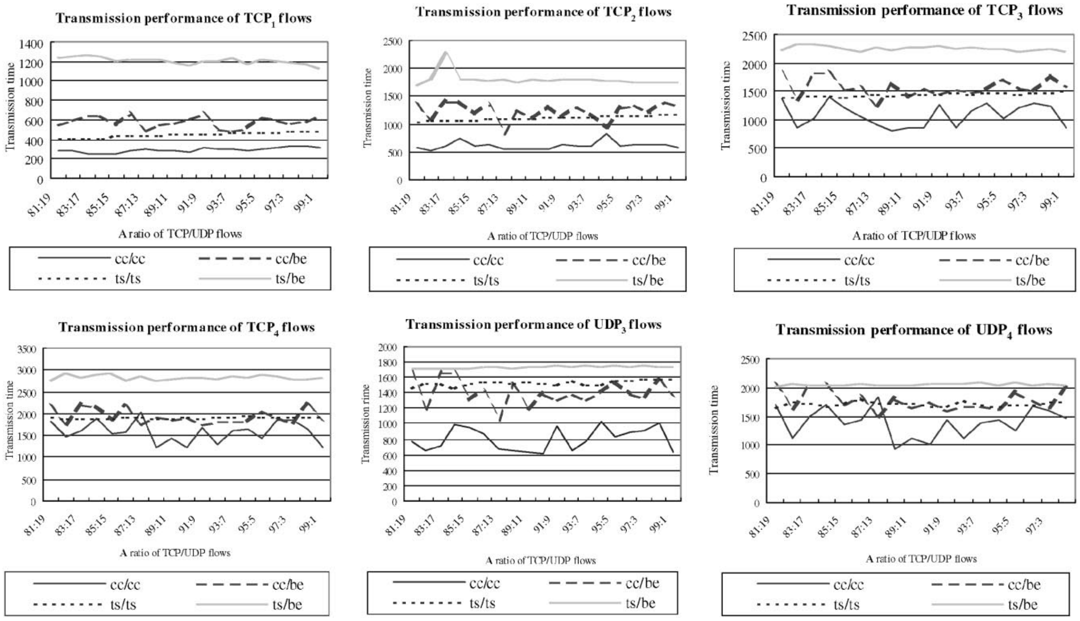

Figure 3. Flows' transmission performance with different ratios of TCP/UDP flows.

Various ratios of TCP/UDP data flows are simulated and the FCFS results are shown in Figure 3. The TCP/UDP ratios do not show significant effects on the traffic performance. According to the different transmission environments, we draw six line charts to 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 a relationship between a ratio of TCP/UDP flows and their transmission performance.

3.5. *Parameter settings of the control mechanisms*

Three important key parameters are investigated: priority ageing, round-robin transmission time, and the number of time slots assigned to each priority. Proper parameter settings allow the control mechanisms to have better control. With numerous simulation settings, one can find that the priority-ageing time is important. The priority-ageing time impacts both mechanisms. Too short a priority-ageing 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. Too long a priority-ageing time, however, may cause a flow with lowest priority to starve because other higher-priority flows may keep coming and jumping ahead of the queue.

The round-robin transmission time is also important in the time-slot mechanism. A flow's round-robin transmission time depends on the number of time slots. A proper number of time slots benefits both TCP and UDP flows. If the number of time slots is too large, a long round-robin transmission time will lead to a first-come-first-served operation. If the number of time slots is too small, the round-robin transmission time can be shorter than the round-trip time. TCP flows cannot get their ACKs from the destination, the retransmission of TCP flows will happen repeatedly and their performance will be poor. Additionally, if differences among round-robin transmission times of different priority flows are too large, this may starve a low priority flow. Otherwise, differential service behaviours do not significantly vary between priority flows. For the time-slot mechanism, the number of time slots assigned to each 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 differential bandwidth allocation of the round-robin transmission time deserves further study.

4. CONCLUSION

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. In these two mechanisms UDP flows are regulated and are not irresponsible to the network congestion.

Because these 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 operational 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. Bernet Y, Blake S, Grossman D, Smith A. An informal management model for Diffserv Routers. *draft-ietf-diffserv-model-04.txt*, July 2000.
2. URL http://www.ietf.org/html.charters/diffserv-charter.html
3. Blake S, Black D, Carlson M, Davies E, Weiss W. An architecture of differentiated services. *RFC 2475*, 1998.
4. URL http://www.vbns.net/stats/flows/html/index.html
5. Tanenbaum AS. *Computer Networks* (3rd edn). Prentice-Hall Inc.: New Jersey, 1996; 413–416.
6. Wright GR, Stevens WR. *TCP/IP Illustrated*, *vol.* 2—*The Implementation*. Addison-Wesley Publishing Company: Reading, MA, 1995; 211–212.
7. McKenney PE. Stochastic fairness queueing. *Internetworking*: *Research and Experience* 1991; **2**:113–131.
8. Floyd S, Van Jacobson. Random early detection gateway for congestion avoidance. *IEEE/ACM Transactions on Networking* 1993; **1**(4):397–413. URL http://www-nrg.ee.lbl.gov/nrg-papers.html
9. Shreedhar M, Lambda GV. Efficient fair queueing using deficit round robin. *Proceedings of the Conference on Applications*, *Technologies*, *Architectures*, *and Protocols for Computer Communication*, September 1995; 231–242.

AUTHORS' BIOGRAPHIES

Chyan Yang received his PhD degree in computer science from the University of Washington, Seattle. Between 1987 and 1992, he worked as an assistant professor in the Department of Electrical and Computer Engineering, U. S. Naval Postgraduate School at Monterey, California. During 1992–1995 he has been with the Institute of Management Science, National Chiao Tung University, Taiwan as an associate professor. He has been with D-Link as a senior advisor between 1995 and 2000. He is currently a Professor in the Institute of Business and Management and the Institute of Information Management, National Chiao Tung University, Taiwan. He is a senior member of IEEE. His research areas cover communication networks and technology management.

Chen-Hua Fu received his BS degree in the Department of Information Management, National Defence Management College, Taiwan, R.O.C., and his MS degree from the Institute of Computer Science, U. S. Naval Postgraduate School in 1993. He is a PhD student in the Institute of Information Management, National Chiao Tung University, Taiwan, R.O.C., where he is working towards his PhD on Network Communication.