Hierarchically aggregated fair queueing for per-flow bandwidth sharing in high-speed backbone networks

Ichinoshin Maki† Hideyuki Shimonishi‡ Masayuki Murata§ Hideo Miyahara†

* Graduate School of Engineering Science,
Osaka University
1-3 Machikaneyama, Toyonaka, Osaka
560-8531, Japan
Phone: +81-6-6850-6588, Fax:
+81-6-6850-6589
E-mail: {i-maki,miyahara}@ics.es.osaka-u.ac.jp

†Comuputer and communication media research, NEC corporation
1-1, Miyazaki 4-Chome, Miyamae-ku,
Kawasaki, Kanagawa, 216-85555, Japan
Phone: +81-44-856-2123, Fax:
+81-44-856-2123
E-mail: simonisi@ccm.cl.nec.co.jp

‡Cybermedia Center, Osaka University
1-30 Machikaneyama, Toyonaka, Osaka
567-0043, Japan
Phone: +81-6-6879-8793, Fax:
+81-6-6879-8794
E-mail: murata@ics.osaka-u.ac.jp

Abstract One of the most important goals for the quality of best effort traffic is fair service among users. Since wide-band access technologies are taking of the limits upon the network resource usage of a single user, fair service is becoming more important service criteria. In this paper, we propose a scalable queue management scheme that realizes per-flow fair service in backbone networks. The scheme allows an arbitrary aggregated queue management according to the forwarding speed of line interfaces; for example, it allows per-flow queue management at relatively slow edge routers but it requires flow aggregation at high-speed backbone routers. To provide per-flow fair service with the flow aggregation, the proposed scheme estimates the number of flows in an aggregated flow and allocates bandwidth in proportional to the estimated number of flows. In addition, to ensure the fairness among flows in the same aggregation, edge routers put a drop preference mark on a packet header so that the flows using relatively large bandwidth are shaped in the core routers. We evaluate the proposed scheme through extensive simulation studies. The results show that the scheme can greatly improve fairness among flows compared to the traditional tail-drop routers. Although it is shown that the fairness is futher improved as the number of flows in one aggregation reduces, high-speed core routers can handle a limited number of aggregated flows and a large number of flows are gathered in one aggregation. Therefore, we also evaluate the tradeoffs between the flow aggregation and fairness. It is concluded that the scheme can offer per-flow fair bandwidth allocation that is comparable to that of per-flow schedulers, even if a small number of aggregated flows that can be handled in core routers is used.