PAPER

Bandwidth Brokers of Instantaneous and Book-Ahead Requests for Differentiated Services Networks

Ying-Dar LIN†, Cheng-Hsien CHANG†, *and* Yu-Ching HSU†, *Nonmembers*

SUMMARY The Quality of Service (QoS) reservations in Differentiated Service (DiffServ) networks can be classified into two sets: Book-ahead (BA) requests and Instantaneous Requests (IRs). When an admitted BA request becomes active, some ongoing IRs is dropped when the bandwidth is insufficient for supporting both IRs and BA requests. The admission control should predict the lifetime, i.e. look-ahead time, of the IRs to prevent the admitted IRs from being dropped. The control should then check whether the available bandwidth during the look-ahead time is sufficient for the incoming IRs. We propose an *application-aware* look-ahead admission control for IRs, which determines the look-ahead time for specific types of IR applications. An admitted BA request might block subsequent ones that could bring more *effective revenue.* Thus, we propose the *deferrable* model of the admission control for BA requests. Simulation results indicate that the application-aware look-ahead admission control successfully reduces the dropping probability and *wasted* revenue of IRs by up to 10 times and 30%, respectively. Besides, the deferrable model indeed results in more BA effective revenue.
key words: *IR, Book-ahead, DiffServ, QoS*

1. Introduction

The Quality of Service (QoS) reservations in Differentiated Services (DiffServ) networks can be classified into two sets, Book-ahead (BA) requests and Instantaneous Requests (IRs). According to the SLA (Service Level Agreement) in DiffServ networks [1], the concept of *Book-ahead* request which books resources by issuing a reservation in advance [2], [3] can be applied in DiffServ networks. Figure 1 illustrates the difference between IRs and BA requests. The BA request carries the *Bandwidth* requirement and the time information, including *Book-ahead* time and *Lifetime*. However, IRs only carries the Bandwidth requirement. The *Arrival Time* is the time when this request arrives. The *Active Time* is the time when this request starts to transmit data. The *Ending Time* is the time when this request terminates. When the BA and IR mechanisms are implemented into bandwidth brokers in DiffServ networks [1], [4]–[6], some services such as Static SLA, Virtual Wire [7], and Tunnels [8] are suitable for BA requests. Note that for simplicity, we focus study the problem in a link.

Three indicators are used herein to evaluate the performance of an admission control. They are *effective revenue*, *IR wasted revenue*, and *IR dropping probability.* The effective revenue is defined as the amount of traffic transferred by complete sessions. The IR wasted revenue is defined as the amount of traffic transferred by dropped IRs. Notably, the revenues to which we refer are normalized. The Independent-Peaks Approximation (IPA) admission control proposed in [9] is not sensitive to application and it drops IRs to pursue higher network utilization. However, higher effective revenue should be pursued instead of network utilization, which equals effective revenue plus IR wasted revenue. A better admission control should achieve higher effective revenue, lower IR wasted revenue and a lower IR dropping probability.

This study addresses three problems. (1) How to adjust BA target utilization to obtain best BA and IR effective revenue. The BA target utilization is defined as the upper bound of the percentage of network capacity that can be allocated to BA requests to prevent IRs from being starved. (2) How to lower the dropping probability of IRs. Figure 2 illustrates the case of ongoing IRs' being dropped. On the Active Time of BA 2, IR 2 and IR 3 are dropped because the residual bandwidth is insufficient for BA 2. To solve this problem, the concept of *look-ahead time* is first presented in [10]. The admission control looks up the *Book-Table* where the booked BA requests are logged, and checks whether the bandwidth is enough for the requesting IR within a *look-ahead time* interval. However, no algorithm is proposed to determine the period of the *look-ahead time.* (3) How to pursue higher BA effective revenue.

Problem (1) is investigated by simulation with varying offered loads of IRs and BA requests. The results prove that the admission control should admit BA requests first and adjust BA target utilization according to BA offered load to obtain higher effective

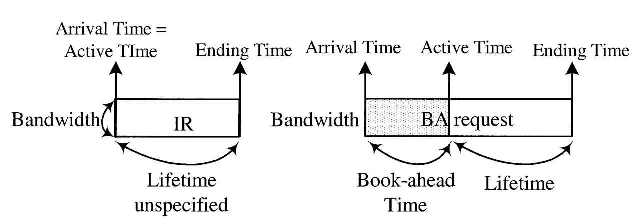

Fig. 1 The difference between IRs and BA requests.

Manuscript received January 4, 2001.
Manuscript revised May 10, 2001.
†The authors are with the Department of Computer and Information Science, National Chiao Tung University, Hsinchu, Taiwan.

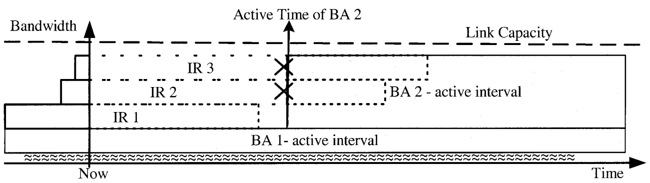

Fig. 2 The case of ongoing IRs being dropped.

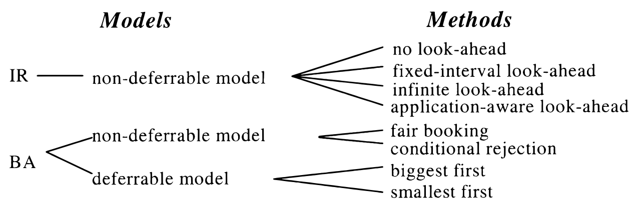

Fig. 3 Classification of BA and IR models and methods.

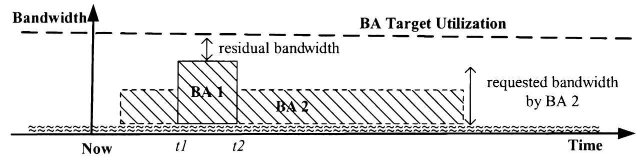

Fig. 4 The case when a BA request blocks the succeeding one which would generate more effective revenue.

revenue. The methods presented to solve problems (2) and (3) are listed in Fig. 3. *Non-deferrable* means that the system should immediately reply accept or reject to the requesting user. For IRs, because the Active Time equals the Arrival Time, only the non-deferrable model exists. The results prove that application-aware look-ahead outperforms the other three simple schemes. For BA requests, the schemes in deferrable model outperform those in non-deferrable model.

The rest of this paper is organized as follows. Section 2 illustrates the proposed methods for admitting BA requests and the data structure of the Book-Table which logs the admitted BA requests. Section 3 presents the methods for IRs. Section 4 illustrates the simulation model and results. Conclusions are finally made in Sect. 5.

2. Operation Models and Methods for BA Requests

To generate more BA effective revenue, a deferrable model is proposed which uses the feature of BA requests that the Arrival Time precedes the Active Time. Notably, this model might be another service class provided by Internet Service Providers (ISPs). However, the period of deferral of the reply is an important issue and discussed in Sect. 4.

The admission control for BA requests can be categorized into two models, non-deferrable and deferrable. In the non-deferrable model, fair booking, the basic scheme for BA requests, involves no heuristic and is compared to the heuristic scheme, *conditional rejection.* In the deferrable model, the admission control defers the response for a certain fraction of the Book-Ahead Time of the incoming BA request. Thus the control may apply heuristic methods on the responses pended BA requests. This work presents two simple and opposite heuristics, *biggest first* and *smallest first.*

2.1 Non-Deferrable Model

Fair booking
Fair booking is the basic and the simplest admission control scheme for BA requests. The information included in a BA request, i.e. Book-Ahead time, Lifetime and Bandwidth, helps the admission control to respond correctly. Fair booking looks up the Book-Table and checks whether the residual bandwidth in the time interval *[Active Time, Ending Time]* is sufficient for the incoming BA request. If so, the request is admitted. Otherwise, the request is rejected.

Conditional rejection
Conditional rejection rejects BA requests that require large bandwidth but have a short lifetime because such requests might cause other requests, which may create more effective revenue, to be blocked. As in the example shown in Fig. 4, when BA 1 is admitted, the residual bandwidth between $t1$ and $t2$ declines causing the succeeding BA 2, which produces more effective revenue than BA 1, to be blocked. However, the above heuristic works only within an interval of offered load. When the load is low, BA requests are seldom blocked. Hence, conditional rejection is disabled till BA blocking rate reaches the threshold, $BABlockProbThreshold$. However, when load is high, the residual bandwidth becomes scarce. If the incoming BA request is rejected, there is little chance for another BA request to utilize the scarce resource. Therefore, when BA effective revenue reaches the threshold, *LatchRatio*, when little bandwidth remains, a latch is added to disable conditional rejection.

2.2 Deferrable Model

Deferring responses to BA requests requires the admission control to track the nearest response time, which is the deadline by which it must reply to a certain request. When the nearest deferred response time is due, the admission control applies one of the following two heuristic methods on the responses pended list. Note that only the requests whose deferred response time is due are replied to.

Biggest first
The BA requests in the response pending list are sorted in descending order of their expected effective revenue, $Bandwidth * Lifetime$. BA requests that generate the

most effective revenue are admitted.

Smallest first

The BA requests in the response pending list are sorted in ascending order of their expected effective revenue. BA requests that generate the smallest effective revenue are admitted.

3. Operation Models and Methods for IR

3.1 Non-Deferrable Model

When an admitted BA request becomes active, some on-going IRs may be dropped because the bandwidth is insufficient for supporting all the IRs as well as the BA request. However, if the lifetime of an incoming IR is known in advance, the dropping can be prevented by simple schemes, such as the fair booking for BA requests. Therefore, the admission control has to *predict* the lifetime, i.e. *look-ahead time*, of incoming IRs. We propose an application-aware look-ahead admission method, which examines the layer-4 header of the IRs and thereby determines the look-ahead time for a specific application type of the IR. Three other simple look-ahead methods—no look-ahead, infinite look-ahead, and fixed-interval look-ahead—are compared with the application-aware look-ahead method.

No look-ahead

The admission control admits an IR if there is enough bandwidth at its Arrival Time. The control does not look up the Book-Table. This scheme is the most aggressive.

Infinite look-ahead

The admission control assumes that the lifetime of incoming IRs is infinite such that the look-ahead time is also infinite. Accordingly, the admission control looks up the whole Book-Table. The IR will be admitted if the residual bandwidth is sufficient for the IR. This scheme is the most conservative.

Fixed-interval look-ahead

In this scheme, the look-ahead time is fixed as the mean lifetime of IRs. The admission control then looks up a part of the Book-Table and checks whether the residual bandwidth is sufficient for the incoming IR.

Application-aware look-ahead

Although the traffic characteristics of flows are different, flows of the same application have similar traffic patterns, such as the amount of transferred traffic per flow (kbits/flow). Further, the type of IR application can be known by examining the layer-4 header. Accordingly, the lifetime can be predicted by *required bandwidth* dividing *the mean of transferred traffic*. This is also the basic concept of the proposed application-aware look-ahead scheme, which determines different look-ahead time for different application type of IRs. Table 1, the outgoing T3 link of National Chiao Tung University [12] monitored by National Center for High-Performance Computing [13], lists the mean amount of traffic transferred of different application types. Notably, this table is used in our simulation.

Table 1 The mean amount of traffic transferred for different application types.

Application type	*Percentage*(%)	*Mean of transferred traffic per IR flow* (kbits/Flow)
WWW	68.3243	49.016
TELNET	4.3817	118.16
ICQ	2.9092	19.336
E-MAIL	2.3708	92.4
FTP	4.2398	3174.32
NEWS	0.7244	643.76
OTHERS	17.0498	441.68
Summary	100	+ 256

4. Simulation Study

4.1 Simulation Model

The IR traffic was monitored in the outgoing T3 link of NCTU campus networks to make our simulations more realistic. Table 1 shows the traffic statistics for IRs. For IRs, the requested bandwidth, b_{IR}, is a Poisson distribution with a mean of 64 kbps. Note that the lifetime of an IR equals x_{IR}/b_{IR} where x_{IR} is the amount of traffic to be transferred. The values of x_{IR} for different application types are determined by Pareto distributions with the shape parameter 1.5 and with the mean values listed in Table 1. The inter-arrival time of IRs is also a Pareto distribution with the same shape parameter 1.5, and with a mean length of $load_{IR}/x_{IR}$ where $load_{IR}$ is the variable load given by IRs.

The mean requested bandwidth for BA requests is a Poisson distribution with a mean of 256 kbps. Furthermore, the lifetime and the book-ahead time are Pareto distributions, where the shape parameter is also 1.5, with a mean length of 1 hour and 1 day, respectively. The inter-arrival time of BA requests is also a Pareto distribution with the same shape parameter 1.5 and a mean length of $load_{BA}/(meanrequestedbandwidth * meanlifetime)$ where $load_{BA}$ is the load given by BA requests which varies in different cases. The network capacity is a 45 Mbps T3 link in our simulation environment.

4.2 Simulation Results

Three metrics used to evaluate the performance are the effective revenue of IR and BA, IR dropping probability, and IR wasted revenue. Note that revenue here is normalized.

Determine BA target utilization

BA target utilization limits the upper bound of the fraction of network capacity that can be allocated to BA requests. It prevents IRs from starvation of bandwidth. The optimal BA target utilization (Fig. 5) under every

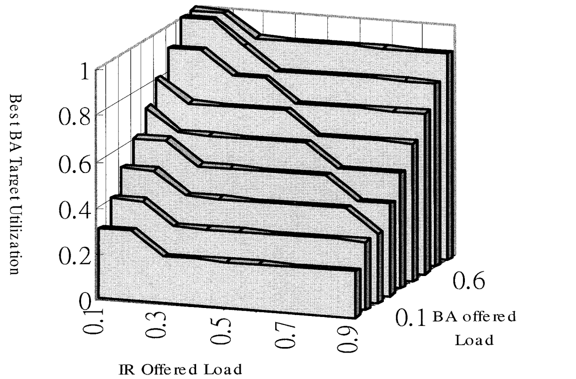

Fig. 5 The optimal BA target utilizations versus different IR and BA offered loads.

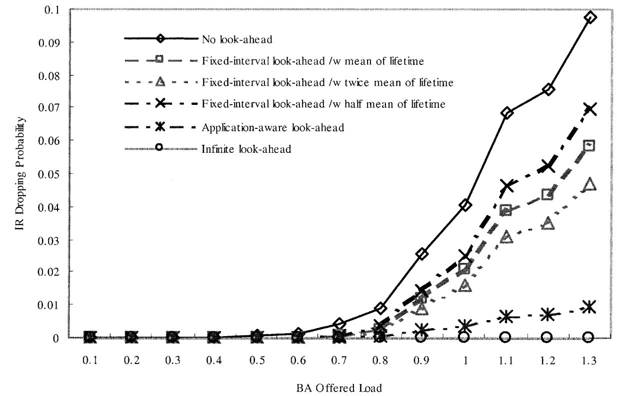

Fig. 6 IR dropping probability versus BA offered load.

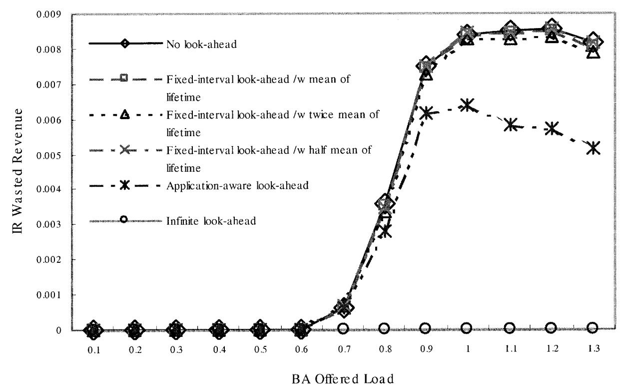

Fig. 7 IR wasted revenue versus BA offered load.

IR and BA offered load combination is determined via simulation. That is, in Fig. 5, various values of target utilization under each combination of IR and BA offered load are simulated and the optimum BA target utilization is obtained. No look-ahead and fair-booking are the respective methods for IR and BA requests to obtain the most conservative result. Figure 5 reveals that the BA target utilization should be set as close as to the offered load of BA requests as possible, because the effective revenue contributed by a BA request is larger than that contributed by an IR. However, when the offered load of BA requests is high, a certain fraction of bandwidth should be reserved for IRs to prevent starvation. Besides, in such a case, IRs is more easily admitted than BA requests.

Comparison between various IR look-ahead methods

This study presents four admission schemes for IRs. Dropping probability, wasted revenue and IR effective revenue are used to evaluate performance. In Figs. 6 to 9, the IR offered load is fixed at 0.7. The BA offered load varies from 0.1 to 1.3. The BA target utilization is 1. The look-ahead time in the fixed-interval scheme is the mean lifetime, half the mean lifetime, and twice the mean lifetime.

The admission control performs look-ahead when admitting an IR to decrease the IR dropping probability and thus reduce the IR wasted revenue. Note that look-ahead time is the predicted lifetime of the IRs. Therefore, from another perspective, the IR dropping probability indicates the precision of the prediction. Figure 6 shows that IR dropping probability of the application-aware look-ahead is up to ten times lower than that of no look-ahead and five times lower than that of fixed-interval look-ahead. When BA offered load increases, the IR dropping probability of the application-aware look-ahead rises more slowly relative to other methods, showing that application-aware look-ahead effectively reduces the IR dropping probability.

Figure 7 shows that the wasted revenues of all the methods are roughly equal except in the cases of application-aware look-ahead and infinite look-ahead. This result follows from the fact that the mean wasted revenues per dropped IR flow, i.e. MWR_{IR}, of other methods are larger. The total IR wasted revenue, WR_{IR}, can be expressed as

$$WR_{IR} = MWR_{IR} * N_{IR} * DP_{IR} \tag{1}$$

where N_{IR} and DP_{IR} are the number of admitted IRs and the IR dropping probability, respectively. Once an IR is dropped, the time that it had been served must be larger than its look-ahead time. For the infinite look-ahead scheme, DP_{IR} is very low and results in the lowest wasted revenue. However, for no look-ahead, DP_{IR} is very high and results in the highest wasted revenue. The lower dropping probability of the application-aware look-ahead, gives less wasted revenue than the others, except the infinite look-ahead.

Figure 7 shows that the IR wasted revenues decreases when the BA offered load is larger than 1.0. This phenomenon, more obvious for application-aware look-ahead because of its lower dropping probability, is due to the higher arrival rate of the BA request leading to a lower value of MWR_{IR}. Figure 8 illustrates the case where the lifetime of the IR is considered infinite. As the arrival rate of BA increases, such that BA requests become active more frequently, the actual life-

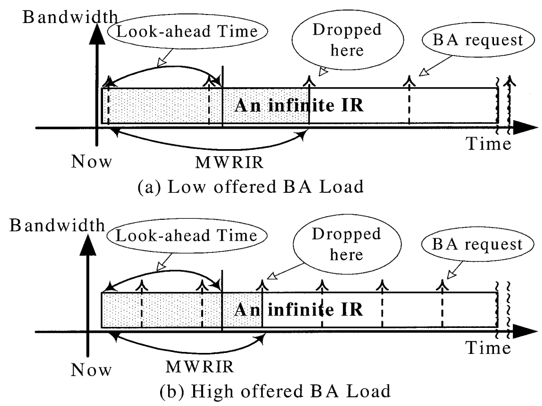

Fig. 8 MWR_{IR} decreases when BA offered load is high.

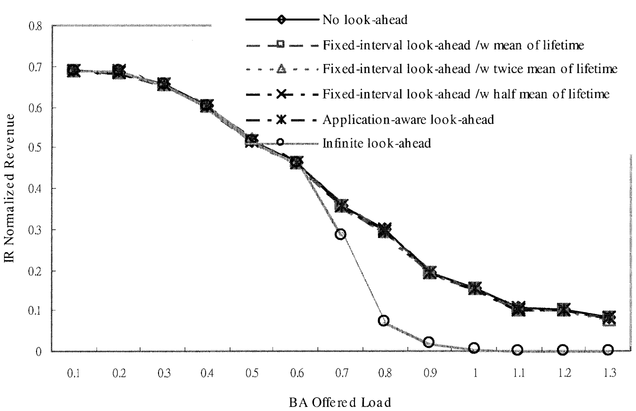

Fig. 9 IR effective revenue versus BA offered load.

time of the IR decreases, showing a smaller MWR_{IR}.

Figure 9 shows the effective revenue of IR. We can see that the curve of the most conservative admission method, infinite look-ahead, decreases rapidly when the BA offered load exceeds 0.6, because the network is then extremely overloaded when the IR offered load, 0.7, is considered. Although application-ware look-ahead exhibits lower dropping probability and lower wasted revenue, the IR effective revenue is roughly equal to that of other methods. Therefore, application-aware look-ahead is a low dropping, low wasted-revenue admission method with no effective revenue degradation.

The effect of the BA conditional rejection method

We use BA effective revenue as an indicator to examine the performance of conditional rejection. In Fig. 10, with no IR traffic offered, the values of the related parameters are as follows. The BA offered load varies from 0.1 to 1.3; BA target utilization is 0.8, *BABlockProbThreshold* is 0.005; *LatchRatio* is 0.9, *BABWThreshold* is 5, and *BALifetimeThreshold* is 0.05. Three methods are examined—fair booking, conditional rejection, and conditional rejection without latch. Notably, in conditional rejection without latch,

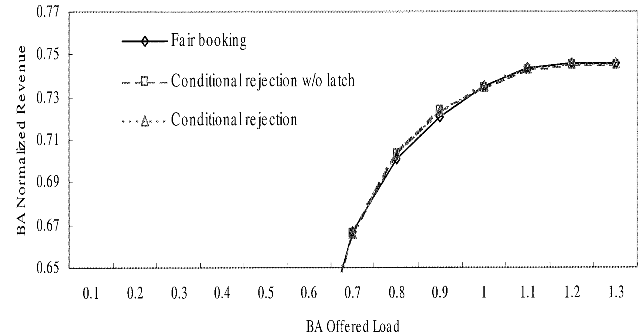

Fig. 10 BA effective revenue versus BA offered load.

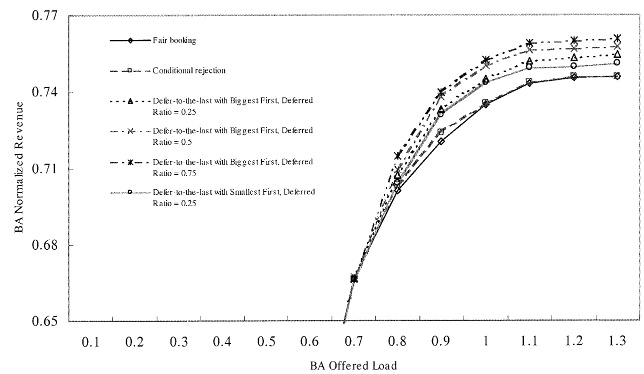

Fig. 11 Comparison of various BA methods.

the *Latch* is always off. As shown in Fig. 10, when BA offered load is between 0.7 and 1.0, the effective revenue of conditional rejection without latch is slightly better than that of fair booking. However, when the effective revenue reaches 90% of the BA target utilization, its rate of increase is reduced. So 90% of BA target utilization is the critical point to enable the latch. However, enabling the latch does not result in significant improvement. Thus, we propose the deferrable model.

Comparison of deferrable BA admission methods

Deferred Ratio is defined as the ratio of the deferring time to the Book-ahead time of BA requests. In Fig. 11, the BA offered load varies from 0.1 to 1.3. The BA target utilization is 0.8 and no IR traffic is offered. According to Fig. 11, the deferrable model performs better than the non-deferrable model. The BA effective revenue increases with the deferred ratio. This observation proves again that the deferrable model out-performs the non-deferrable model because the admission control can better decide when more information is provided. In fact, the opposite algorithm, smallest-first, though worse than biggest-first, is still much better than other methods of the non-deferrable model. Accordingly, the best way to increase the BA effective revenue is to change the fundamental operation model.

5. Conclusion

This work presents several methods to solve (1) the BA target utilization adjustment problem, (2) the IR dropping problem, and (3) the BA effective revenue improvement problem. Simulation results have shown that the BA target utilization should be adjusted according to the BA offered load to acquire higher effective revenue, but is somewhat insensitive to the IR offered load. Our results further indicated that application-aware look-ahead exhibits an IR dropping probability up to ten times less than that of no look-ahead and five times less than that of fixed-interval look-ahead. Besides, application-aware look-ahead can reduce IR wasted revenue while maintaining the same level of effective revenue. For BA requests, the deferrable model leads to better effective revenue than the non-deferrable model.

Future effort could be as follows. An optimal admission decision algorithm could be developed. Techniques in mathematical programming could be used to model and optimize the problem. An optimal algorithm may be impractical due to the lack of availability of some data or computational complexity but could act as an upper bound for reference.

References

[1] S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang, and W. Weiss, "An architecture for differentiated services," RFC 2475, Dec. 1998.
[2] O. Schelen and S. Pink, "An agent-based architecture for advance reservations," 22nd Annual Conference on Local Computer Networks, 1997.
[3] D. Ferrari and A. Gupta, "Admission control for advance reserved real-time connections," Third IEEE Workshop on the Architecture and Implementation of High Performance Communication Subsystem, 1995.
[4] J. Heinanen, F. Baker, W. Weiss, and J. Wroclawski, "Assured forwarding PHB group," RFC 2597, June 1999.
[5] V. Jacobson, K. Nichols, and K. Poduri, "An expedited forwarding PHB," RFC 2598, June 1999.
[6] K. Nichols, S. Blake, F. Baker, and D. Black, "Definition of the differentiated services field (DS Field) in the IPv4 and IPv6 headers," RFC 2474, Dec. 1998.
[7] V. Jacobson, K. Nichols, and K. Poduri, "The 'Virtual wire' behavior aggregate," draft-ietf-diffserv-ba-vw-00.txt, March 2000.
[8] D. Black, "Differentiated services and tunnels," draft-ietf-diffserv-tunnels-00.txt, Feb. 2000.
[9] A.G. Greenberg, R. Srikant, and W. Whitt, "Resource sharing for book-ahead and instantaneous-request calls," IEEE/ACM Trans. Networking, vol.7, pp.10–22, Feb. 1999.
[10] O. Schelen and S. Pink, "Sharing resources through advance reservation agents," IWQoS'97, 1997.
[11] O. Schelen, A. Nilsson, J. Norrgard, and S. Pink, "Performance of QoS agents for provisioning network resources," IWQoS'99, 1999.
[12] http://www.nctu.edu.tw/
[13] http://moat.nchc.gov.tw/moat/

Ying-Dar Lin was born in Hsi-Lo, Taiwan in 1965. He received the Bachelor's degree in Computer Science and Information Engineering from National Taiwan University in 1988, and the M.S. and Ph.D. degrees in Computer Science from the University of California, Los Angeles in 1990 and 1993, respectively. At UCLA Computer Science Department, he worked as a Research Assistant from 1989 to 1993 and worked as a Teaching Assistant from 1991 to 1992. In the summers of 1987 and 1991, he was a technical staff member in IBM Taiwan and Bell Communications Research, respectively. He joined the faculty of the Department of Computer and Information Science at National Chiao Tung University in August 1993 and is Professor since 1999. His research interests include design, analysis, and implementation of network protocols and algorithms, wire-speed switching and routing, quality of services, and intranet servers. He has been a consultant for several high-tech companies and authored two books including a textbook on Computer Network Experiments. Dr. Lin is a member of ACM and IEEE. He can be reached at ydlin@cis.nctu.edu.tw and http://www.cis.nctu.edu.tw/˜ydlin.

Cheng-Hsien Chang was born in Taiwan in 1976. He received the B.S. and M.S. degrees in Computer Science from National Chiao Tung University, HsinChu, Taiwan in 1998 and 2000, respectively. His research interests include design, analysis, and implementation of network protocols and algorithms, wire-speed switching and routing, quality of services, and intranet services.

Yu-Ching Hsu was born in Tei-Tou, Taiwan in 1973. She received her B.A. and M.S. degrees in Computer and Information Science from National Chiao Tung University in 1995 and 1997, respectively. She is currently a Ph.D. candidate in the same department. Her research interests include protocol design, analysis, and wireless communications. She can be reached at ychsu@cis.nctu.edu.tw.