Two-Stage Dynamic Uplink Channel and Slot Assignment for GPRS

Yu-Ching HSU[†], Ying-Dar LIN[†], *and* Mei-Yan CHIANG[†], *Nonmembers*

SUMMARY   General packet radio service (GPRS) uses a two-stage mechanism to allocate uplink radio resource to mobile stations (MSs). In stage-1, the base station (BS) assigns several packet data channels (PDCHs) to an MS. Furthermore, a PDCH may be assigned to multiple MSs. In stage-2, therefore, the BS selects one of the multiplexed MSs in a PDCH to use the radio resource. In this paper, maintaining a load balance between PD-CHs in stage-1 is examined and several selection schemes to lower the mis-selection rate in stage-2 are proposed. From our simulation results, the cost deduced from the poor load balancing and selection schemes render a lower system throughput and a non-negligible increase in packet queuing delay. Among the various stage-2 selection policies, round robin with linearly-accumulated adjustment (RRLAA) has the lowest mis-selection rate and outperforms the one without any heuristic by up to 50%.
key words: *GPRS, channel assignment, slot assignment, scheduling*

1.   Introduction

GPRS [1], developed by European Telecommunications Standard Institute (ETSI), is one of the standards of Global System for Mobile communications (GSM) Phase 2+. To accommodate the bursty behavior of Internet traffic, GPRS is designed as a packet switching system. To support various QoS requirements of MSs, GPRS supports multislot assignment, which enables an MS to transmit data on several PDCHs in parallel. Besides, GPRS multiplexes several MSs in a PDCH, i.e. overbooking, to quickly accommodate the resource assignment to the traffic pattern of the uplink flows and thus utilize uplink radio resource to its full advantage. To distinguish the multiplexed MSs in a PDCH, Uplink State Flag (USF) is employed to number the MSs. Notably, only one of the multiplexed MSs can send packets on this PDCH each time. These uplink resource allocation mechanism specified in GPRS [1] is viewed herein as a *two-stage* assignment. When a data channel request is received, the BS performs the stage-1 assignment to assign PDCHs to the requesting MS. That is, it sends the resource assignment message containing a list of PDCHs and their corresponding USFs to the MS. Therefore, by indicating the corresponding USF on the downlink as a stage-2 assignment, the BS assigns a time

slot to the MS that can transmit on this PDCH in the next time slot.

To improve GPRS uplink resource utilization and throughput, three primary directions have been investigated. The first regulates the random access channels shared by voice and data traffic [2-5]. The second meliorates the Automatic Repeat reQuest/Forward Error Correction (ARQ/FEC) mechanisms [6],[7]. The third assesses the capacity of GPRS system [8]. Much research has been devoted to improving GPRS uplink radio resource utilization. However, little has been published on optimizing the dynamic allocation procedure, that is, considering both multislot assignment and multiplexing.

The objective of this study is to examine the effect of various load balancing principles in stage-1 and different selection schemes in stage-2. In stage-1, two PDCH load measurement methods, Number of Assigned Flow (NoAF) and Effective Transmission over Last Cycle (EToLC), are compared with the simple one, Fixed Number of PDCH's (FNoP). In stage-2, because multple MSs might be multiplexed in a PDCH and only one of them can transmit at a time, to utilize the precious uplink radio resource, the BS has to *predict* who has data to send and then assign the following time slot to that MS. Therefore, prediction, or selection, policies for stage-2 assignment are also proposed herein.

In stage-2, several selection policies, each composed of the concepts of round robin, linearly-accumulated or exponentially-accumulated adjustments, are proposed and compared. As Round Robin with Linearly-Accumulated Adjustment (RRLAA) has the lowest mis-selection rate of these proposed schemes with about 5% difference, only its simulation results are presented. Both system throughput and mean packet queueing delay of RRLAA are compared with those of OPTimal (OPT) and Pure Round Robin (PRR). The simulated results have demonstrated that EToLC and RRLAA perform better within stage-1 and stage-2 assignments, respectively. Notably, in our focus on two-stage dynamic uplink channel and slot assignments for GPRS data traffic, voice requests are not considered.
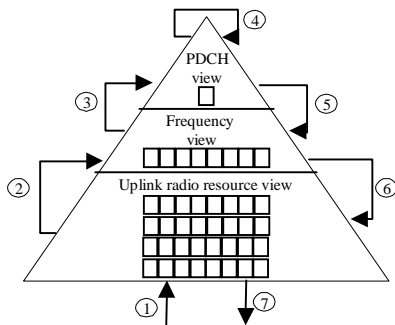
This paper is organized as follows. Section 2 describes the proposed method for both stage-1 and stage-2 assignments. Section 3 describes the simulation model and numerical results in which the proposed two-stage assignment schemes are analyzed and compared.

1. BS receive PCR (Packet Channel Request) from the MS.
2. Decide the number of PDCHs to be assigned to the MS according to the QoS parameters in the PCR. Assume that K PDCHs are requested here. ($0 < K <= 8$)
3. Select K PDCHs with minimum assigned load in each frequency.
4. For those selected PDCHs, examine whether the number of multiplexed MSs will exceed the multiplexing threshold if one more MS is added.
5. Compute total load of selected PDCHs in each frequency that none of the selected PDCHs will exceed the multiplexing threshold if one more MS is added.
6. Select the frequency with minimum total load.
7. If none of these frequencies meet, then block occurs. Reply block notification to the MS. Otherwise, reply PRA (Packet Resource Assignment) which indicates the assigned frequency, PDCHs, and the corresponding USFs to the MS.

Fig. 1 Procedure of stage-1 assignment.

Finally, section 4 contains a discussion regarding future work.

2. Two-Stage Dynamic Channel and Slot Assignment

After the BS receives the Packet Channel Request from an MS, it assigns uplink radio resource to the MS in two stages. In stage-1, multiple PDCHs with the corresponding USFs are assigned to an MS to meet the QoS requirement. Note that a PDCH may be assigned to multiple MSs. The criterion of the stage-1 assignment procedure designed herein balances the PDCHs loads.

Though a PDCH may be assigned to multiple MSs, only one MS is selected to transmit data at a time. Therefore, in stage-2, to utilize the radio resource, the BS has to predict who has data to send and then assign the following time slot to that MS. Therefore, prediction policies for stage-2 assignment are also proposed herein.

2.1 Stage-1 Channel Assignment

After receiving the Packet Channel Request, besides the number of PDCHs, BS must decide the specific PDCHs to be assigned to the MS. The number, which should not exceed 8, can be decided based on the QoS parameter of the Packet Channel Request, but deciding which PDCHs to assign is more critical. Notably, as USF is encoded by 3 bits, the number of MSs assigned to a particular PDCH should not exceed the multiplexing threshold, 8. Fig. 1 illustrates the PDCHs-selection procedure of stage-1.

When the normalized load of some PDCHs ex-

ceeds 1, while others are lower than 1, limited system throughput is predicted. That is because those exceeding 1 will result in a maximum normalized PDCH utilization of 1 and others are lower than 1, i.e. not all offered load is pumped. However, if the loads between PDCHs are balanced and are lowered than 1, all the offered load will be pumped. In addition, when the load is high, the packet queuing delay also tends to be high. This phenomenon spurs an observance of the differences in system throughput and average packet queuing delay when the PDCH load-balancing scheme differs. Two load measurement methods are proposed to balance PDCHs loads and thus, utilize the radio resource completely. They are NoAF and EToLC.

Notably, the term "flow" indicates the traffic generated by an MS. Because in stage-1 a PDCH cannot be assigned to an MS twice, the number of multiplexed MSs within a PDCH is the same as the number of assigned flows. In the proposed two methods, the behavior of the flows in a PDCH, e.g. whether a packet is transmitted, is used to measure the load of that PDCH in different ways. That is different load metrics are adopted in the two methods.

Number of assigned flow (NoAF)

In this case, the number of multiplexed MSs within a PDCH is chosen as the load measurement metric without considering either traffic characteristics or the actual behavior of each flow. If K PDCHs are requested, BS first locates K PDCHs with minimal assigned MSs in each frequency. Second, BS assigns the MS to the K minimally loaded PDCHs in the frequency with the lowest total load of K minimally loaded PDCHs. Without considering the actual traffic behavior, this scheme can be considered as frequency-wise and PDCH-wise balanced.

Effective transmission over last cycle (EToLC)

Supposing M MSs are multiplexed in a PDCH, a PRR (Pure Round-Robin) cycle is then defined as each of the M MSs has an opportunity to transmit once within a cycle. Thus, the length of a PRR cycle becomes M. The load metric employed by EToLC is defined as the number of transmissions occurred during the *previous* PRR cycle. Then, the BS assigns the MS PDCHs in a manner similar to that described in Fig. 1. The factors of measured load thus include not only the number of assigned flows in a PDCH and the actual behavior of each flow, but also the accuracy of stage-2 selection policy. Notably, the definition of "cycle" in different stage-2 assignment policies might differ.

2.2 Stage-2 Slot Assignment

As multiple MSs are multiplexed in a PDCH and only one of them can transmit at a time, the BS has to predict who has data to send and then assign the following timeslot to that MS. If the selected MS has no data impending, the slot is wasted. Therefore, an ac-

curate prediction scheme will promote the uplink radio resource utilization.

To quickly accommodate the bursty characteristics of Internet traffic and thus increase the success rate in selecting an MS to transmit data, a reward and penalty policy is applied to the stage-2 assignment scheme. That is, the selected MS should be rewarded if the assigned timeslot is employed, otherwise a penalty should be issued. The reward is an improved chance of being selected, while the penalty is a lesser chance.

In this paper, two stage-2 assignment schemes, PRR and RRLAA, are proposed and their detailed procedures are described below. In addition, an OPT assignment scheme is also presented and compared with the above two assignment schemes.

Pure Round-Robin (PRR)

In PRR, each multiplexed MS in a PDCH is round-robined to use the uplink channel. Without considering actual behavior of each flow, all MSs are assumed having impending data to send. Therefore, the selection sequence is not affected by whether an MS utilized the last assigned timeslot or not. Furthermore, the length of a PRR cycle equals the number of MSs multiplexed in this PDCH.

Round-Robin with Linearly-Accumulated Adjustment (RRLAA)

RRLAA is derived from PRR and contains the concepts of penalty and reward. Its basic principle is to reduce the transmission chance for the MSs that failed to utilize the last assigned slot, and increase the chance for those who had. For RRLAA, a *Penalty* cycle and a *Reward* cycle are defined and appear alternately. A Penalty cycle is derived from a PRR cycle by skipping MSs who waste their last assigned timeslots in Penalty cycles. The more times the MS wastes the assigned timeslots, the stiffer penalty deserved, i.e. the times of being skipped increases linearly. Thus, an MS will be skipped in n successive Penalty cycles when it wastes n successive assigned timeslots in Penalty cycles. However, when the MS begins to send packets, the penalty accumulation is reset and becomes zero. Hence, the length of a Penalty cycle is less than or equal to that of a PRR cycle.

To execute the reward policy, *Reward* cycles is defined and is inserted between Penalty cycles. An MS is authorized to transmit during the following Reward cycle if it transmits data in the previous Penalty cycle. Furthermore, the number of admitted transmission in a Reward cycle also increases linearly. That is an MS will be rewarded n timeslots in a Reward cycle when it successively employs the assigned timeslots in n Penalty cycles. Furthermore, within a Reward cycle, the conceptual sequence of assignment is also round robin. Note that, an MS will be selected to send data at most once in a Penalty cycle but possibly multiple times in a Reward cycle. Besides, the reward and the penalty for each MS are originally zero. The detailed
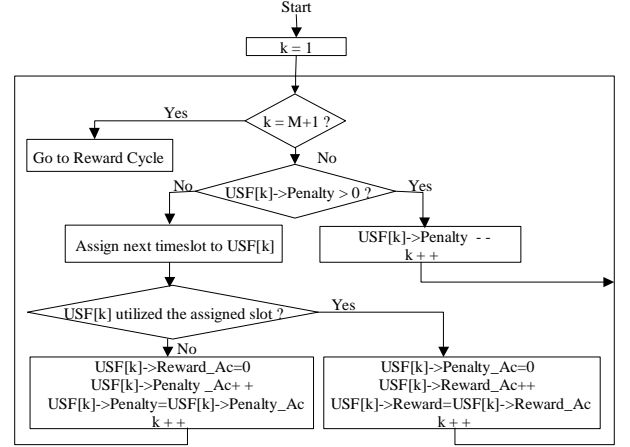

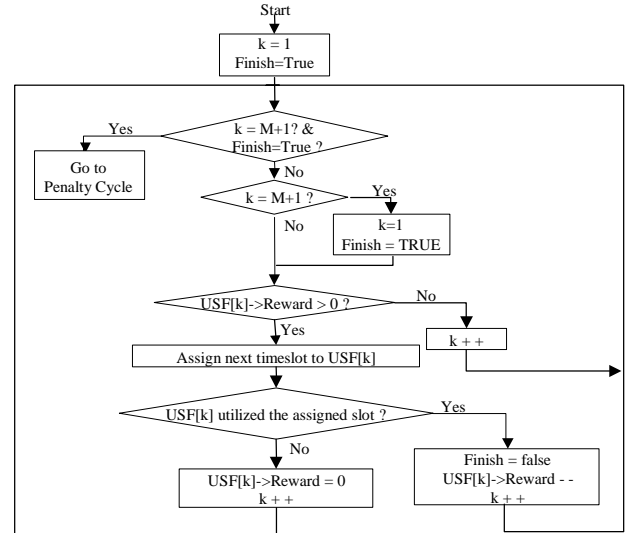**Fig. 2** The procedure to process a Penalty cycle.

Fig. 3 The procedure to process a Reward cycle.

procedures to process a Penalty cycle and a Reward cycle are depicted in Fig. 2 and 3, respectively. Notably, it is assumed that M MSs are multiplexed in a PDCH.

OPTimal (OPT)

To compare the performance of PRR and RRLAA, OPT is introduced. OPT is based on the assumption that whether an MS has data to send or not is known in advance. OPT adheres to the basic PRR scheme by not selecting the MSs who have no impending data. Thus, only when none of the multiplexed MSs have data to transmit, is the uplink radio resource wasted. Notably, OPT does not contain the reward and the penalty policies. Fig. 4 presents the detailed procedure.

Fig. 4   Procedure of OPT.

3.   Results and Discussion

3.1   Simulation Model

The numerical results have been obtained via an event-driven simulator, PARSEC [9]. Furthermore, 32 PD-CHs are assumed for uplink data traffic. In addition, to simulate the traffic behavior of each MS, a two-level ON/OFF model is applied. In the first level, i.e. connection level, the ON and OFF periods are exponentially distributed and their mean lengths are 0.64 and 1.024 seconds, respectively. When each MS enters the connection ON state, the number of PDCHs assigned to it is randomly decided and ranges between 1 and 8. In the second level, i.e. the packet level, the packet interarrival time is modeled by the Pareto distribution, where the shape parameter is 1.7. Furthermore, the mean packet interarrival time is 0.00722125 seconds. Notably, the traffic generator of the second level is enabled only when the MS is in connection ON state. Besides, the performance is mainly affected by $offered\ load$, which is aggregated by connection level and packet level traffic. In this paper, the offered load is adjusted by the number of MSs.

The following metrics evaluate the performance of load balancing and selection schemes for stage-1 and stage-2 assignments:
System Throughput: the ratio of effective transmission over simulation time.
Standard Deviation of PDCH Utilization: the standard deviation of the utilization of all PDCHs.
Standard Deviation of the loads contributed to the assigned PDCHs by an MS: The standard deviation in the view from an MS.
Improvement: the improved ratio in system throughput of the proposed scheme to that of a simple scheme, i.e. RND-RND operation model described in the next subsection.
Mis-selection rate: the rate of erroneously choosing an MS which has no impending data when at least one active MS multiplexed in the same PDCH has.
Average Packet Queueing Delay: the average time interval between a packet is generated and transmitted.

3.2   Simulation Results

Comparison between Load Balancing Schemes for Stage-1 Channel Assignment
    NoAF and EToLC are proposed for stage-1 assignment. To examine their effects, four operation models are simulated, RND-NoAF, RND-EToLC, RND-RND, and FNoP-NoAF. Table 1 lists the comparison among these operation models. For RND-RND, RND-NoAF and RND-EToLC, the number of assigned PDCHs are randomly decided. By contrast, for FNoP-NoAF, the number is fixed and the same for all MSs. Further,
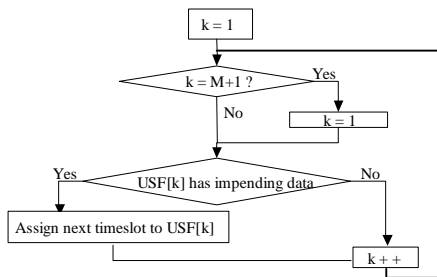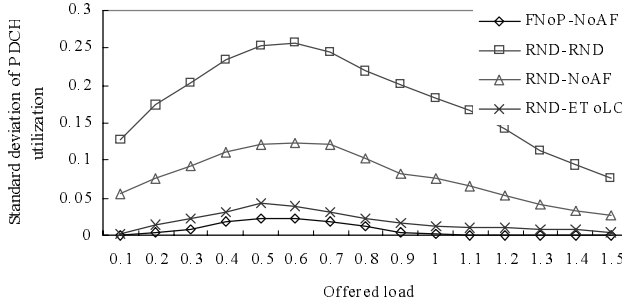
Table 1   Assumptions of different load balancing schemes.

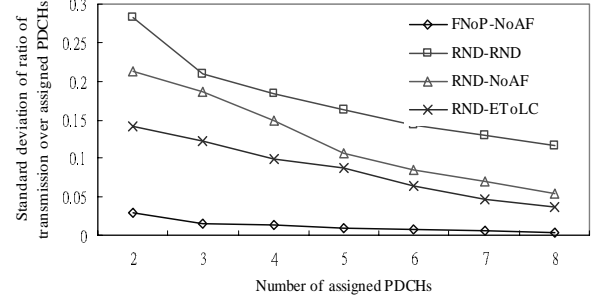| Scheme	Number of assigned PDCHs	PDCH load metric
FNoP-NoAF	Fixed	Number of assigned flows
RND-RND	Random from 1 to 8	No metric applied
RND-NoAF	Random from 1 to 8	Number of assigned flows
RND-EToLC	Random from 1 to 8	Effective transmission over last cycle

Fig. 5   Standard deviation of PDCH utilization for different stage-1 operation models.

Fig. 6   Standard deviation of the loads contributed to the assigned PDCHs by an MS.

Fig. 7   System throughput for different load balancing schemes.

to avoid a decline in system performance that is due to poor selection by the stage-2 assignment, OPT is adopted. Notably, RND-RND simulates the basic case when no heuristic is applied. Further, FNoP-NoAF is designed to be compared with other three models and is considered as the most load-balanced case. This is mainly because that the MSs have the same traffic pattern and the number of assigned PDCHs. Therefore, within FNoP-NoAF, MSs contribute the same traffic loads to the assigned PDCHs. Furthermore, since NoAF scheme is applied, the number of MSs assigned to PDCHs is balanced. Hence, FNoP-NoAF is the most load-balanced case. However, it is not the case in real life because MSs will have different QoS requirement.

The most instinctive method to determine whether the offered load is balanced among PDCHs or not is to observe the standard deviation of PDCH utilization. If the loads of PDCHs are balanced the standard deviation of PDCH utilization should be low, and vice versa. Fig. 5 reveals that, excluding FNoP-NoAF, RND-EToLC scheme has the lowest standard deviation when MSs have been assigned different numbers of PDCHs, which is also the normal case as each MS has a different QoS requirement. Observing the curves in Fig. 5 are lower when the offered load is either getting higher or lower. That is because when total load is low, all PDCHs are light-loaded; thus, less of a deviation range exists, and vice versa when total load is getting high.

Fig 6 shows the standard deviation of the loads contributed to the assigned PDCHs by an MS. For the same reason, if the load is quite balanced among

PDCHs and stage-2 assignment is OPT, a flow shall distribute the same load to the assigned PDCHs and the deviation should be low. The same observation, i.e. EToLC scheme outperforms NoAF, is also obtained from the figure.

Fig. 7 depicts the system throughput for the four operation models. Expectedly, the measured system throughput of FNoP-NoAF, i.e.the most load-balanced model, is nearly equal to the offered load. However, due to a better PDCH load balancing, RND-EToLC outperforms RND-NoAF. This is because, in RND-NoAF, the normalized offered loads of some PDCHs exceed 1, while others do not. However, not all the offered load can be pumped in the over-loaded PDCHs. Therefore, when the overall offered load is high, the measured throughput does not reach the offered load.

Fig. 8 depicts the improvement of the proposed RND-NoAF and RND-EtoLC as compared to the RND-RND scheme. When offered load is near 1, the improvement rate is expressly high. This is because that, in RND-RND operation model, many PDCHs are over-loaded and others are not; thus, not all offered load can be pumped. Hence, poorer system throughput than that of either RND-EToLC or RND-NoAF is resulted in.

Fig. 9 and 10 illustrate the performance of FNoP-NoAF operation model with the varying fixed number of assigned PDCHs to an MS. Notably, the stage-2 assignment is PRR. In the figures, $M = 1$ indicates that each MS is assigned a PDCH, and each PDCH is multi-
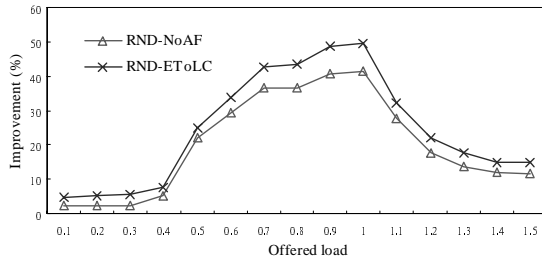
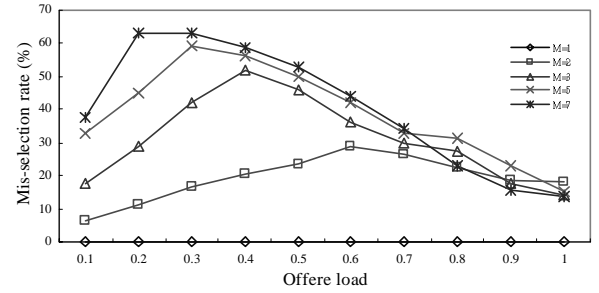**Fig. 9** Mis-selection rates for different fixed number of assigned PDCHs to an MS.
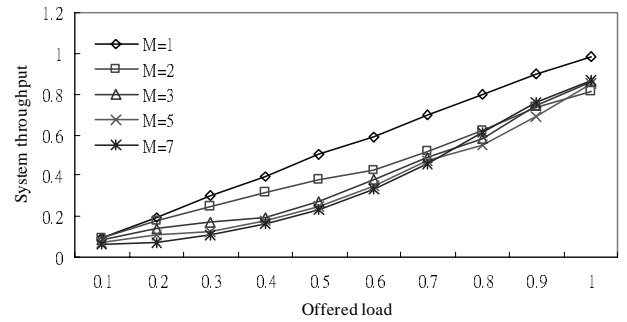
Fig. 10 System throughput for different fixed number of assigned PDCHs to an MS.

Fig. 8 Improvement in system throughput for proposed load-balancing schemes over the RND-RND scheme.

plexed with only one MS. Furthermore, $M=2$ indicates that each MS is assigned two PDCHs, and each PDCH is multiplexed with two flows, and so on.

When each PDCH is assigned to only one flow, mis-selection does not occur. Therefore, single slot assignment combined with no multiplexing in PDCHs outperforms in mis-selection rate and system throughput (Fig. 9 and 10). Moreover, when offered load is low, the fewer flows multiplexed in a PDCH, the lower the mis-selection rate, and vice versa. This is interesting that assign too many PDCHs to a low-loaded flow will increase the flow's chance of being selected and thus increases mis-selection rate. However, for a high-loaded flow, the larger number of assigned PDCHs will increase the flow's chance to send data. Hence, more PDCHs should be assigned to high-loaded flows, and vice versa for low-loaded flows.

Comparison between Selection Schemes for Stage-2 Slot assignment

Our study for stage-2 assignment focuses on the performance of selection schemes. Through the simulation, the mis-selection rate of each selection scheme is compared and, thus discover that the mis-selection should receive penalties. Except OPT and PRR, all possible combinations of round robin, with either exponential or linear penalty or reward, are simulated. Since RRLAA outperforms other combinations with about 5% difference, only the results of OPT, PRR and RRLAA are shown. Notably, the stage-1 assign-
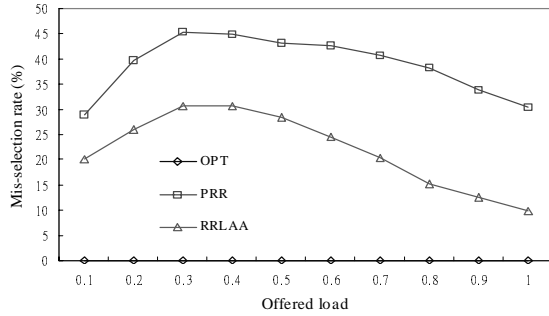
Fig. 11 Mis-selection rate of different selection schemes.

Fig. 13 Average packet queueing delay for different selection schemes.
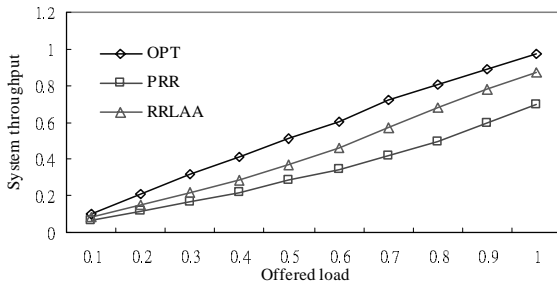

Fig. 12 System throughput for different selection schemes.

ment used is RND-EToLC.

Fig. 11 shows the mis-selection rates of different stage-2 schemes. The rate of OPT is always zero because of the assumption that whether the flow has impending data or not is known in advance. The figure also illustrates that RRLAA has lower mis-selection rate than PRR due to the reward and penalty policies, which consider the actual behavior of each flow. The other observation is that when offered load is either extremely low or high, the mis-selection rate is lower. This is because that when the offered load is low, the probability of all MSs multiplexed in a PDCH having no impending data is high. Besides, when the offered load is high, the probability of having data for transmission for all MSs assigned to a PDCH is also high. Therefore, the rates are lower when the offered load is either extremely low or high.

Fig. 12 shows the system throughput for different selection schemes. According to mis-selection rates, the system throughput of OPT is better than RRLAA, which is better than PRR. Furthermore, the lower the mis-selection rate is, the higher the system throughput. This is because *system throughput* can be approximately expressed as $offered\ load * (1 - mis\text{-}selection\ rate)$.

Another penalty of mis-selection is longer packet queuing delay because once a timeslot is wasted, the queueing time of all the impending packets becomes longer. Fig. 13 verifies that the average packet queuing delay of PRR is almost three times as long as that of RRLAA. In addition, the queuing delay of both PRR
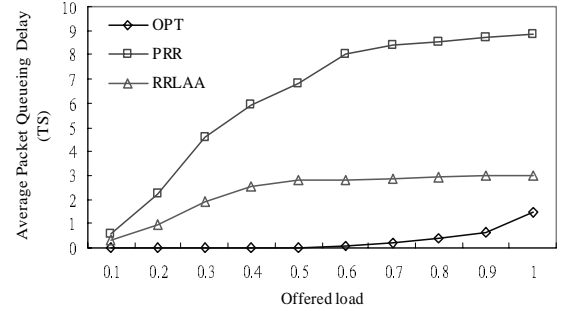
and RRLAA converges as the offered load increases. It is because that, evidently, the average packet queuing delay is proportional to both the mis-selection rate and the offered load. When the offered load of increases, the mis-selection rate decreases (Fig. 11). Therefore, when the offered load increases, both curves converge. However, when offered load exceeds 1, the queue length will grow infinitely and thus results in infinite packet queuing delay. Notably, for OPT, the delay is only proportional to the offered load; thus, the curve increases with the offered load.

4. Conclusion

Two GPRS PDCH load-balancing schemes for stage-1 channel assignment and one selection scheme for stage-2 slot assignment, have been proposed herein. To maintain load balancing between PDCHs, two load metrics, NoAF and EtoLC, were proposed and compared. When the load of a PDCH is measured, NoAF considers only the number of assigned flows, whereas EToLC considers both the number of assigned flows and the actual behavior of each flow. For stage-2 assignment, a selection scheme, called RRLAA, is presented. RRLAA contains the concept of linearly accumulated reward and penalty. Reward means assigning more timeslots to a flow and occurs when the chosen flow utilized the timeslot, and vice versa for penalty.

From this, several conclusions can be drawn. For stage-1 assignment, considering the actual behavior of each assigned flow within a PDCH helps maintaining load balancing between PDCHs. That is EToLC outperforms NoAF. Among the various stage-2 selection policies, Round Robin with Linearly-Accumulated Adjustment (RRLAA) has the lowest mis-selection rate and outperforms the one without any heuristic up to 50%. Moreover, selection schemes influence both system throughput and packet queuing delay.

References

[1] ETSI, "GSM 03.64 Overall description of the GPRS radio interface, stage 2," v7.0.0, Release 1998.

[2] Shaoji Ni, Hilggman,S., "GPRS performance estimation in GSM circuit switched services and GRPS shared resource systems", *IEEE WCNC*, 1999.

[3] Shaoji Ni, Yong Liang, Sven-Gustav Haggman, "Outage Probability for GPRS over GSM Voice Services", *IEEE WCNC*, 1999.

[4] J.E. Wieselthier and A. Ephremides, "Fixed- and Movable Boundary Channel-Access Schemes for Integrated Voice/Data Network," *IEEE Transaction on Communication*, Vol. 43, No. 1, pp. 64-74, January 1995.

[5] S. Ghani and m. Schwartz, " A Decomposition Approximation for the Analysis of Voice/ Data Integration," *IEEE Transaction on Communication*, Vol. 42, No. 7, pp. 2441-2451, July 1994.

[6] Meyer. M, "TCP performance over GPRS", *IEEE WCNC*, 1999.

[7] Wessam Ajib, Philippe Godlewski, "Acknowledgment procedures at radio link control level in GPRS", *Proceedings of the 2nd ACM international workshop on Modeling, analysis and simulation of wireless and mobile systems*, pp. 33-40, 1999.

[8] Calin, D., Malik, S., Zeghlache, D., "Traffic scheduling and fairness for GPRS air interface" *IEEE VTC*, Fall, 1999.

[9] http://pcl.cs.ucla.edu/

on Computer Network Experiments. Dr. Lin is a member of ACM and IEEE. He is the founder and director of Network Benchmarking Lab (NBL) which reviews the functionality, performance, conformance, and interoperability of networking products. He can be reached at ydlin@cis.nctu.edu.tw and http://www.cis.nctu.edu.tw/ ydlin.
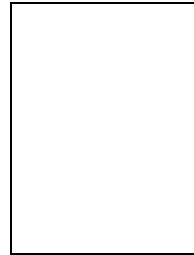
Mei-Yan Chiang   Mei-Yan Chiang was born in Tai-Chung, Taiwan in 1976. She received her Bachelor's degree in Computer Science and Information Engineering from National Chiao Tung University in 1998 and the M.S degree in Computer and Information Science from National Chiao Tung University in 2000. She has been a software engineer in EDA industry since August 2000. She can be reached at mychiang@springsoft.com.tw.

Yu-Ching Hsu   Yu-Ching Hsu was born in Tei-Tou, Taiwan in 1973. She received her M.S. and Ph.D. in Computer and Information Science from National Chiao Tung University in 1997 and 2002, respectively. She has been a researcher in Industrial Technology Research Institute since August 2002. Her research interests include protocol design, analysis, and wireless networks, 3G networks. She can be reached at YuChing@itri.org.tw.
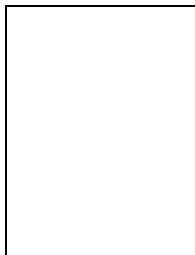
Ying-Dar Lin   Ying-Dar Lin was born in Hsi-Lo, Taiwan, in 1965. He received the Bachelor's degree in Computer Science and Information Engineering from National Taiwan University in 1988, and the M.S. and Ph.D. degrees in Computer Science from the University of California, Los Angeles in 1990 and 1993, 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, network security, and content networking. He has been a consultant for several high-tech companies and authored two books including a textbook