

Figure 1: GPRS Network Architecture (Courtesy: <http://communications.siemens.com/repository/31/3129>)

need to be transferred for all inter-BTS handoffs when the PCU is co-located with the BTS, while it needs to be transferred only for inter-BSC handoffs when the PCU is co-located with the BSC.

An implementation of GPRS/EDGE typically has the PCU co-located with the BSC. The PCU terminates the RLC protocol and transports the useful bits along with the headers to the BTS, which then performs the channel coding before transmitting the block over the air to the intended mobile. In this paper, we restrict attention to those implementations that terminate the RLC protocol at either the BSC or the SGSN, while doing the channel coding down at the BTS. We refer to the link between where the RLC protocol is terminated and where the channel coding is performed as the backhaul link. This backhaul link can be either circuit switched or packet switched.

A circuit backhaul for GPRS/EDGE would mean that we provision and reserve bandwidth in the backhaul for each of the air interface timeslots separately. Typical circuit backhaul links are T1 or E1 leased lines, and one timeslot in the T1/E1 link is called a DS0, whose bandwidth is 64 kbps. In Table 1, we show the possible coding schemes of GPRS and the possible modulation and coding schemes (MCS) of EDGE and their corresponding data rates. In order to allow sufficient bandwidth in the backhaul for MCS 9, which corresponds to an over-the-air throughput of 59.2 kbps, we have to provision one DS0 (64 kbps) for each data-capable air interface timeslot. For circuit backhaul, regardless of whether or not there is a mobile on an air interface timeslot, one DS0 needs to be allocated to the air interface timeslot.

Table 1
Coding Schemes and Data Rates for GPRS and EDGE

GPRS Data Rates			EDGE Data Rates		
Coding Scheme	Modulation	Data Rate/Timeslot (kbps)	Modulation and Coding Scheme	Modulation	Data Rate/Timeslot (kbps)
CS-1	GMSK	8	MCS-1	GMSK	8.8
CS-2		12	MCS-2		11.2
CS-3		14.4	MCS-3		14.8
CS-4		20	MCS-4	17.6	
			MCS-5	8-PSK	22.4
			MCS-6		29.6
			MCS-7		44.8
			MCS-8		54.4
			MCS-9		59.2

Instead of nailed circuit connections, if we packetize the backhaul we could provision bandwidth for all the data-capable air interface timeslots together. Packetized backhaul allows statistical multiplexing gains thereby enabling backhaul capacity reduction. The statistical multiplexing benefits

3 CARRIER SELECTION ALGORITHMS

3.1 Throughput Estimation for Circuit Backhaul Carriers

When a mobile requests a new connection setup, the first step is to determine the carrier in which the mobile is to be scheduled. We can accomplish this by estimating the throughputs that will be offered to the new mobile by each of the carriers (if the mobile is scheduled on that carrier) and then selecting the carrier that gives the highest throughput estimate. We obtain the throughput estimate by determining the average number of opportunities per block period that a mobile will get and taking the product with the throughput per data timeslot for the new mobile on that carrier. For a circuit backhaul carrier i , an estimate of the throughput can be expressed as follows: $R_i = T_i \bar{R}$, where R_i is the estimated throughput for the user in carrier i , T_i is the estimated number of opportunities that the user will get per block period, and \bar{R} is the average data rate that the user will get if the user gets one opportunity per block period. The quantity \bar{R} depends on the average channel condition of the user and can be based on past estimates on the average data rates users get. All these quantities depend on the user's characteristics, such as, whether the mobile is a GPRS or an EDGE mobile, and the mobile's slot capabilities.

The quantity R_i depends on several variables as described below.

- It depends on the existing load on the carrier, which in turn depends on the number of users already scheduled on carrier i and their current coding schemes. As one can expect, the higher the load on the carrier, the lower the throughput for the new user.
- It depends on the number of idle timeslots on carrier i on which no mobile is allocated before the arrival of the new user.
- It depends on the total number of data timeslots on the carrier.
- It also depends on the average MCS that will be employed by the new mobile on the circuit backhaul carrier. This can be estimated based on the type of the new mobile (GPRS/EDGE) and any other prior knowledge such as the channel conditions for the incoming user.
- It depends on the multislot capability of the mobile, i.e., the number of timeslots on which the new mobile can operate in a single block period.
- It also depends on the QoS requirement of the new mobile.
- It depends on the scheduler, which arbitrates and decides which among the mobiles on a given timeslot gets served in a given block period. There are several scheduler choices, both channel unaware [4], and channel aware [5], [6].

Note that for the simple case of a round-robin scheduler, when all the timeslots are occupied, one can approximate T_i as the ratio of the number of timeslots in the carrier to the total number of mobiles on the carrier. Based on the estimate of the throughputs on different carriers R_i , we choose the best carrier i^* from the set \mathcal{C} of all carriers as $i^* = \arg \max_{\mathcal{C}} R_i$

3.2 Throughput Estimation for Packet Backhaul Carriers

The equation for estimating throughputs for packet backhaul carriers needs to be modified because no specific backhaul can be attributed to each data timeslot of a packet backhaul carrier as opposed to circuit backhaul carriers where each data timeslot always has one DS0 capacity associated with it. We propose three different methods to obtain the throughput estimate of a packet backhaul carrier. All the three methods have the general framework where the throughput estimate made assuming circuit backhaul carrier is scaled down by a factor γ_i . Thus all the three methods use the following general expression, but use different mechanisms to estimate the value of γ_i : $R_i = \gamma_i T_i \bar{R}$. We denote by N_i the number of data timeslots in carrier i , and by B_i the backhaul bandwidth provisioned for the data timeslots in carrier i in kbps.

The scaling factor γ_i is given by

$$\gamma_i = \min \left\{ 1, \frac{B_i}{U_i} \right\}, \quad (2)$$

where U_i is the average throughput of the carrier after adding the new user assuming an unconstrained backhaul. The quantity U_i is obtained as follows:

$$U_i = \sum_{j \in S_i} T_i^j R_j + T_i \bar{R} \quad (3)$$

where S_i is the set of all users in carrier i excluding the new user, and R_j is the data rate per transmission opportunity per block period at the current coding scheme of user j (as given by Table 1). The quantity T_i^j is the number of transmission opportunities that user j will get in carrier i after the new user is added. We compute an estimate of the average throughput per data timeslot including the new mobile for the packet backhaul carrier assuming there were no backhaul restrictions. We use the current coding schemes for all existing mobiles and the estimated MCS for the new user to determine the per timeslot throughput U_i .

It is evident from the above equations that we do not scale down unless $U_i > B_i$, i.e., unless the load is sufficiently high that the aggregate throughput exceeds the backhaul bandwidth. On the other hand, the static scaling scheme always reduces the throughput as long as the carrier backhaul capacity is smaller than the maximum backhaul required for an EDGE mobile, and the enhanced scaling scheme reduces the throughput by just as much once all the data timeslots are occupied.

This solution is more accurate but slightly more complex than the static scaling and enhanced static scaling options. The complexity arises from the computation of the average throughput per data timeslot. This scheme does not distinguish between two carriers with the same average throughput estimate per data timeslot but different carrier backhaul capacities so long as the average throughput estimate per data timeslot is smaller than carrier backhaul bandwidths. Ideally, we should prefer the carrier with higher backhaul so that even if the mobiles start operating at higher coding schemes, there is still sufficient backhaul bandwidth. These can be overcome by appropriately breaking ties in favor of carriers with higher backhaul.

Once a carrier has been chosen, the exact time slot configuration is allocated to the mobile based on the capability of the mobile and the loads existing on the individual timeslots of the carrier. We do not deal with this problem in this paper.

We carried out simulation-based analysis of the three different options, where we compare the cell throughput, fairness, and complexity performance of the different options. See Section 5 for a partial set of results. Due to space constraints, we have not included all the results. Based on these results, we conclude that the dynamic scaling option, although more complex, is superior in performance compared to the other options. The results also show that the enhanced static scaling option is better than static scaling option under lightly loaded situations.

4 SCHEDULING ALGORITHMS

In this section, we first describe a naive approach to account for the reduced backhaul, and then describe four different options that variously trade off between fairness, efficiency, and complexity of the algorithms. The four options also consider the different alternatives of reducing MCS, searching for alternate mobiles whose MCSs are lower, and working with updated or un-updated sort metrics. We use the ‘‘class and channel condition based weighted proportional fair scheduler’’ described in [1] as the baseline original scheduling algorithm. Each mobile has an associated sort metric, which captures the mobile’s channel condition, its QoS requirements, and the service it has received so far. The sort metric of a mobile gets updated when the mobile is scheduled on a time-slot. The update rule takes into account the current channel condition of the mobile and the fairness-throughput trade-off factor. In each time-slot in a block period, the contending active mobiles (i.e.,

4.2.3 Option 3: Search for alternate mobiles until schedule is just above backhaul

The third approach is the same as the second approach, except that we do not accept the replacement if it takes the bits scheduled below the backhaul limit. Instead, we stop replacing mobiles when the next replacement results in the scheduled bits falling below the backhaul limit, and go with option 1 from that point onwards. This is done because the MCS reduction option gives us fine granular reduction in scheduled bits, which the option of replacing mobiles may not be able to give.

4.2.4 Option 4: Reduce MCS based on pre-updated sort metrics

The fourth approach is designed to reduce the processing requirements for computing a schedule that fits within the backhaul. Here we do not update the sort metrics, but just pick the mobile with the largest pre-updated sort metric, and reduce the mobile's MCS until either the schedule fits within the backhaul capacity, or the mobile is eliminated from the schedule. If the mobile gets eliminated from the schedule, we pick the mobile with the next highest pre-updated sort metric and repeat the process. At the cost of a slight increase in unfairness, we will be able to reduce the processing requirements significantly using this approach. Note that the reduction in complexity is coming not only from not updating sort metrics to begin with, but also from the fact that we do not have to recompute the mobile with the largest sort metric after each MCS reduction.

5 NUMERICAL RESULTS AND DISCUSSION

In this section, we present simulation results to validate the different carrier selection and scheduling algorithms. We used a GPRS/EDGE system simulator for our simulation results. This system simulator has faithful implementations of the RLC/MAC and physical layers, and has support for different traffic types. The metrics of interest are the overall average cell throughput and the average user throughput. For the carrier selection, we simulated a system with 3 carriers all having packetized backhaul. The backhaul for the individual carriers were 3 DS0s, 4 DS0s, and 5 DS0s, respectively, and all the timeslots on all the carriers are dedicated data timeslots. All mobiles belong to class 4 which means that they can operate in 3 timeslots in a single block period in the downlink and 1 timeslot in the uplink. All the mobiles were assumed to have the same quality of service requirements. Each timeslot can have up to 4 mobiles allocated to it, and one of them will be scheduled during each block period.

We consider two different scenarios - heavy loading and light loading. In the heavily loaded scenario, all the timeslots have four mobiles assigned to them during the entire simulation period and each mobile has enough data to transmit over the entire duration. We study the performance of the algorithms when carriers are deployed under different reuse patterns. Figure 2 shows the average cell throughput and the average user throughput achieved in the downlink with different EDGE fractions under three different reuse patterns. An EDGE fraction of 0.25 implies that 25 % of the mobiles are EDGE mobiles and the rest are GPRS mobiles. In the figures, SS, ESS, and DS denote the static scaling, enhanced static scaling, and dynamic scaling solutions, respectively.

As the figure shows, the overall cell throughput is consistently higher for dynamic scaling option and this improved performance of dynamic scaling option is more prominent at more conservative reuse patterns (e.g. 4x3). As mentioned earlier, the ESS option performs exactly the same as SS option when all timeslots have at least one mobile allocated to them, which is the case for the entire duration in a heavily loaded system. Hence we are not able to observe difference in performance between SS and ESS schemes. All the above observations also hold for the average per user throughput. As expected, the overall cell throughput and the average per user throughput values are higher for conservative reuse patterns where interference is low.

Under the lightly loaded scenario, we design a system where the input source rate is 100 kbps. Hence, in this situation, there are possibilities of idle timeslots and timeslots with fewer than four mobiles allocated on them. Figure 3 compares the performance of the three schemes in terms of the average throughput obtained by a user for different EDGE fractions. In this case also, we observe that DS option offers slightly better performance than the other two schemes in most of the

6 CONCLUSIONS

In this paper, we designed and analyzed scheduling and carrier selection algorithms for a packetized backhaul and showed the benefits from making these algorithms aware of the backhaul available for each carrier. We considered three different carrier selection algorithms that variously make use of total backhaul bandwidth, number of occupied timeslots, and current MCSs of mobiles to estimate the throughput the new user will receive in each carrier. We showed that the dynamic scaling approach performs better than the other two options. We designed four different scheduling algorithms for taking the backhaul bandwidth into account. We showed that these four options perform better than the naive scheduler option. We also found that the four options offer similar performance. Therefore, we recommended option 4, which has significantly less computational complexity.

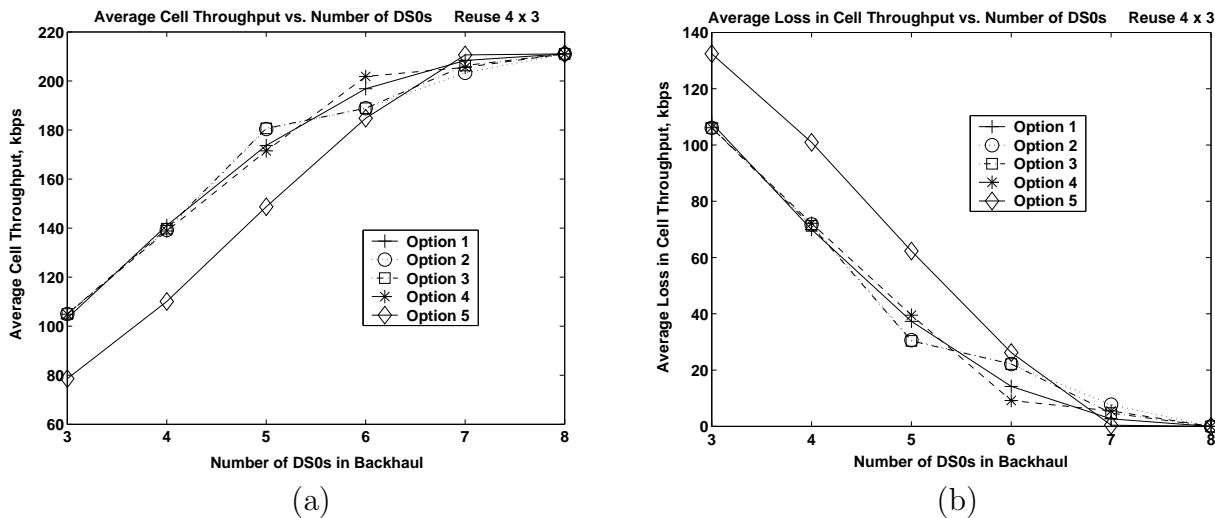


Figure 4: Performance comparison of the scheduling algorithms (a) Average cell throughput (kbps) (b) Average cell throughput loss (kbps)

References

- [1] R. Agrawal, A. Bedekar, R. J. La, R. Pazhyannur, and V. Subramanian, "Class and channel condition-based scheduler for EDGE/GPRS" *Proceedings of SPIE ITCOM*, pp. 59–69, Aug. 2001.
- [2] 3rd Generation Partnership Project, 3GPP TS 04.60, Release 1999, "General Packet Radio Service (GPRS); Mobile Station-Base station system interface; Radio link Control/Medium Access Control (RLC/MAC) protocol" Available from www.3gpp.org.
- [3] 3rd Generation Partnership Project, 3GPP TS 23.060, Release 1999, "General Packet Radio Service (GPRS); Service description; Stage 2" Available from www.3gpp.org.
- [4] T. Halonen, J. Romero, and J. Melero, "GSM, GPRS, and EDGE Performance: Evolution towards 3G/UMTS" 2nd edition, John Wiley and Sons, 2003
- [5] X. Liu, E. K. P. Chong, and N. B. Shroff, "A framework for opportunistic scheduling in wireless networks" *Computer Networks*, vol. 41, no.4, pp. 451–474, March 2003.
- [6] S. Borst, "User-level performance of channel-aware scheduling algorithms in wireless data networks" *Proceedings of IEEE Infocom 2003*, San Francisco, USA.