Revenue based Call Admission Control for Wireless Cellular Networks

Srihari Nelakuditi Raja R Harinath Sanjai Rayadurgam Zhi-Li Zhang
Department of Computer Science & Engineering
University of Minnesota, Minneapolis, MN 55455
{srihari,harinath,rsanjai,zhzhang}@cs.umn.edu

ABSTRACT

Call admission control (CAC) schemes in wireless cellular networks attempt to reduce call dropping probability possibly at the expense of increased call blocking probability. We propose using channel reassignments in a controlled manner to minimize call dropping while maintaining high spectrum utilization. Guard channels are used to control the number of reassignments. The number of guard channels is dynamically determined using reassignment frequency as feedback. A simple scheme that attempts to maintain the number of reassignments under a specified target is described. A revenue based CAC scheme is then presented which attempts to maximize income by balancing the penalty for reassignments against the reward for serviced calls. Simulation results confirm and validate the ideas discussed in this paper.

I. INTRODUCTION

A mobile cellular system consists of cells, each of which has a base station (BS) that serves mobile hosts (MH) in that cell. To establish a communication session, a MH requests the BS in its cell for a channel. BS assigns an unused channel if available. Otherwise the call is blocked. Channel assignment schemes suggested in the literature [4] range from fixed to dynamic, centralized to distributed, and timid to aggressive. Aggressive algorithms reassign channels being used by calls in progress to accommodate new requests. If a MH moves from one cell to another during a session, the responsibility of continuation of service is handed over to the new BS. This is known as *handoff*. If the new BS cannot find an idle channel, the on-going session must be terminated. Since forced termination of an existing call is less desirable than blocking a new call, call admission control schemes [3], [5], [6], [8] have been proposed to reduce the call dropping probability, possibly at the expense of increasing call blocking probability.

Most CAC schemes are based on the guard channel concept [3]. This approach offers a generic means of improving the probability of successful handoffs by simply reserving a few channels in each cell exclusively for handoffs. This may, however, result in low spectrum utilization. A recently proposed scheme, based on the concept of shadow cluster [5], performs CAC decisions using estimates of future resource requirements of mobiles. Besides its complexity, this scheme requires knowledge about the mobility pattern of users, which may not always be available.

In this paper, we explore the use of channel *reassignments* to reduce handoff failures. It has been pointed out that dynamic channel assignment schemes cannot maximize channel reuse as they serve randomly offered call attempts. This leaves enough room for channel reconfigurations which can be used to service otherwise unsatisfiable handoffs. Additionally, a handoff results in releasing a channel in the old cell which can potentially be used for reassignment to free up a channel in the new cell. But reassignments may inconvenience users and incur communication and processing overheads. However, at low loads the additional number of calls serviced more than compensates the overhead due to reassignments. As the system gets overloaded, the number of reassignments grows rapidly without a corresponding increase in the number of serviced calls. Hence reassignments must be prudently used to be beneficial.

We propose using guard channels to control reassignments. Reassignment in the neighborhood is used as an indication of congestion by a cell. The number of guard channels in a cell is dynamically adjusted based on the reassignment frequency of its neighbors. This significantly reduces the need for further reassignments by having the cell react prior to onset of congestion, while maintaining high spectrum utilization. It is possible to conceive various adaptive approaches that utilize the knowledge of reassignments in the neighborhood of a cell. A simple scheme would be to keep the reassignments under a given target. An ideal, but perhaps impractical, scheme is one that maximizes the revenue, given that each call serviced brings in some reward and each reassignment incurs some penalty. We propose a scheme that attempts to balance the number of reassignments and the number of admitted calls by dynamically adjusting guard channels. Simulation results show that this scheme uses reassignments profitably, while achieving near zero dropping probability.

The rest of the paper is organized as follows. Section II describes the simulation set-up that is used for all the experimental results reported in this paper. Section III discusses the use of guard channels to reduce call dropping and their effect on bandwidth utilization. We then show in Section IV, how reassignments can be used to almost eliminate handoff failures and investigate the effect of load and mobility on reassignments. A simple scheme that attempts to keep the number of reassignments under a specified target is described. We then present a revenue based scheme. We conclude the paper with some remarks about future extensions to this work.

II. EXPERIMENTAL SETUP

All our experiments were conducted using a simulation model described here. The cellular system consists of 144 hexagonal cells arranged as a 12x12 array. To avoid the boundary effect, this array is wrapped-around in both the dimensions. The size of a *cluster*, the set of neighboring cells that cause mutual co-channel interference, is assumed to be 7, amounting to a total of 18 interfering cells surrounding each cell. This interference zone is referred to as the cell's *neighborhood*.

All simulations were carried out assuming that the total number of channels in the system is 70. Local Packing [2], a distributed dynamic channel allocation algorithm, is used to allocate channels to new call or handoff requests. Each cell maintains an augmented channel occupancy (ACO) table that tracks channels used and the

0-7803-4912-1/99/$10.00©1999 IEEE

Cell-site	Channel Number								Num Free
	1	2	3	4	5	6	...	M	
i		X			X		...		0
i_1	X		X				...		0
i_2						X	...	X	2
i_3	X		X				...		1
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮
i_{k_i}			X				...	X	3

Fig. 1. ACO table at cellsite i

number of free channels available in each neighboring cell. Figure 1, shows an example ACO table at cell i. Channels not in use in a cell's neighborhood are eligible for assignment to calls in that cell. When no such channel is available, a channel that is being used in exactly one neighboring cell having free channels is considered eligible for reassignment. For example, given the ACO table of cell i in Figure 1, the mobile currently using channel 6 in neighbor i_2 can be relocated to one of the 2 free channels available in i_2, releasing channel 6 to satisfy a request in cell i.

Call interarrival time and call holding time are assumed to be exponentially distributed. The mean call holding time, $1/\mu$, is fixed at 1 and thus the average load per cell λ/μ is varied by varying the mean arrival rate per cell, λ. The number of handoffs per call is assumed to be geometrically distributed with mean $\bar{h}$ and hence the probability of a call moving out of a cell is $\bar{h}/(\bar{h}+1)$. Mobiles are assumed to migrate to any of the adjacent cells with equal probability. A mobile resides for the same amount of time at each visit to a cell during the lifetime of the call. Except when studying the effect of varying mobility, $\bar{h}$ was set to 3 for all the experiments. Similarly the load is fixed at 6 Erlangs/cell when studying the effect of mobility. Each experiment simulated 100,000 calls.

III. UTILITY OF GUARD CHANNELS

Admission control is based on the principle that denial of service to new calls is better than unreliability of service to admitted calls. In other words, dropping calls in progress due to handoff failures is less desirable than blocking new calls. A simple way of giving priority to handoffs is to reserve a number of channels (say g) exclusively for servicing handoffs. These reserved channels are called *guard channels*. A cell rejects new calls whenever the number of free channels available in that cell goes below g. While this reduces the chances of handoff failures, it may result in under-utilization of scarce bandwidth.

Unlike fixed channel allocation schemes, with dynamic schemes like LP guard channels are not exclusive to a cell. A single unused channel in a neighborhood may be counted as a free channel by more than one cell. Hence, g guard channels per cell does not imply that two neighboring cells have a combined total of $2 \cdot g$ channels reserved for handoffs. In fact, this could be as low as g when both cells have the same set of free channels available. A new call is admitted only if the number of free channels available is more than g. In this paper, GUARD refers to a guard channel based scheme using LP for channel allocation.

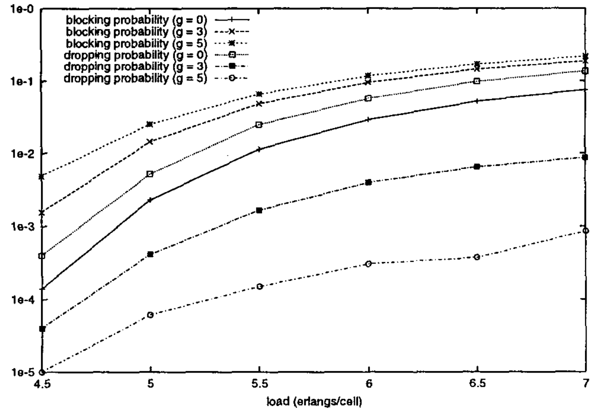

Fig. 2. Effect of guard channels on blocking and dropping probabilities

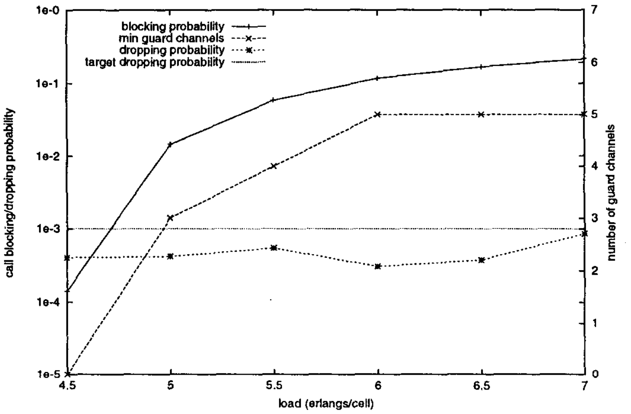

Fig. 3. Minimum number of guard channels required to keep dropping probability below $1/1000$ for varying load and the corresponding blocking probabilities

Figure 2 shows the blocking and dropping probabilities under varying load for different values of g. As the number of guard channels increases the dropping probability decreases while the blocking probability increases. When g is zero, the dropping probability is even more than the blocking probability. This is due to the fact that there are more handoffs ($\bar{h}$ is 3) than new calls and handoffs are not given any preferential treatment over new calls. As the load increases both blocking and dropping probabilities increase across all values of g. A similar trend can be seen (not shown here) as mobility increases.

The objective of admission control schemes is to minimize the call blocking probability while keeping call dropping probability below an acceptable limit. For guard channels based schemes, this is equivalent to guaranteeing a given upper bound on call dropping probability using as few guard channels as possible. However, statically determining the right number of guard channels to achieve this is not possible, since it depends on traffic conditions such as load and mobility.

Figure 3 shows the minimum number of guard channels per cell required to keep call dropping probability below $1/1000$, for different values of load. It also shows the corresponding block-

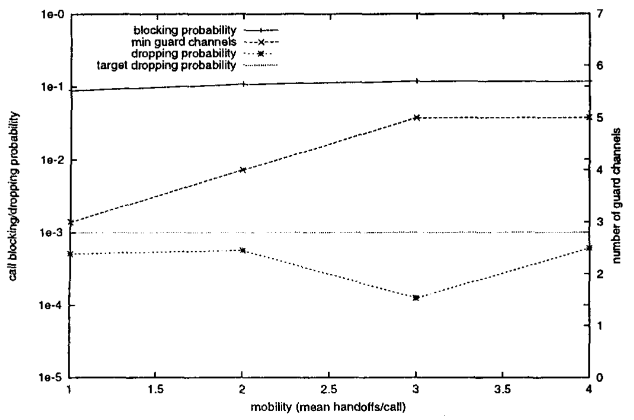

Fig. 4. Minimum number of guard channels required to keep dropping probability below 1/1000 for varying mobility and the corresponding blocking probabilities

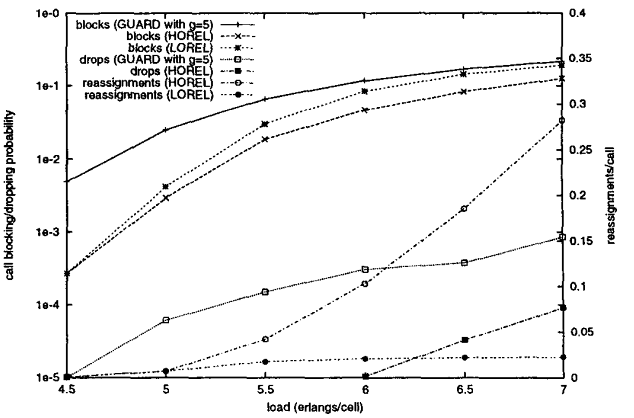

Fig. 5. Illustration of the effectiveness of reassignments

ing probabilities. These values were obtained for each load by conducting a series of experiments with progressively increasing numbers of guard channels until the dropping probability fell below 1/1000. The number of guard channels is shown on the right vertical axis. It can be seen that the minimum number of guard channels required varies with load. For example, at a load of 5 Erlangs/cell the number of guard channels required is 3 while it is 5 when load is 6 Erlangs/cell.

Figure 4 shows the minimum number of guard channels per cell required to keep call dropping probability below 1/1000 for different values of mobility. Again, the minimum number of guard channels required varies with mobility. For example, when mobility ($\bar{h}$) is 2, the number of guard channels required is 4 while it is 5 when $\bar{h}$ is 3.

From these figures, it is evident that fixing the number of guard channels statically will result in either under-utilization of bandwidth or poor quality of service. Hence it is desirable to have an adaptive scheme that dynamically adjusts the number of guard channels in each cell to suit the prevailing traffic conditions. In the following section, we show how reassignments could be used for this purpose.

IV. EFFECTIVENESS OF REASSIGNMENTS

When no free channels are available in a cell for a new request, it may be possible to relocate an on-going call in a neighboring cell to a different unused channel in that cell and release the currently used channel. This released channel can then be assigned to the new request. Such a reallocation of channels to calls in progress to make room for a new request is called *reassignment.* This has been shown to increase the capacity of the system to carry more load [1]. But reassignments may inconvenience users and incur communication and processing overheads. Hence reassignments should be used sparingly and judiciously.

We propose using reassignments to reduce handoff failures. Dynamic channel assignment schemes cannot maximize channel reuse as they serve randomly offered channel requests. As a result, it is almost always possible to reassign channels to calls in progress to make room for a successful handoff. Further, a handoff involves releasing a channel in the old cell which can potentially be reassigned to free up a channel in the new cell. While reassignments do have an associated cost they are much more preferable than dropping calls. Thus, relying on reassignments to avoid call dropping is justifiable. However, it is not advisable to use reassignments to admit a new call since the need for reassignment indicates overload in the neighborhood. Admitting a new call in such a scenario is likely to cause many more reassignments. Moreover, it may lead the system to a state of saturation where no further reassignments are possible resulting in failed handoffs.

A simple scheme based on reassignments, called HOREL, rejects a new call only if there are no free channels available (i.e., no guard channels) but use reassignments if needed to satisfy a handoff request. Figure 5 compares the performance of GUARD with g of 5 and HOREL. The average number of reassignments per call is shown on the right vertical axis. This indicates the expected number of times a mobile is forced to switch channels due to reassignments. HOREL drops less than 1 in 10000 calls while blocking much fewer calls than GUARD. But as the load increases HOREL causes quite a few reassignments with more than 1 reassignment per every 4 calls admitted, at load 7. At low loads the additional number of calls serviced more than compensates the overhead due to reassignments. As the system gets overloaded, the number of reassignments grows rapidly without a corresponding increase in the number of serviced calls.

It is clear that the uncontrolled use of reassignments even for handoffs could be counter-productive. However, used in a controlled manner, apart from ensuring success of handoffs, reassignments also provide an indication of congestion in the neighborhood. One way to control reassignments is by employing guard channels. Reserving a few channels for handoffs obviates the need for frequent reassignments. On the other hand, reassignments can be thought of as a feedback mechanism to determine the right number of guard channels to eliminate handoff failures. The number of guard channels can be dynamically adjusted based on the extent of congestion as indicated by the reassignment frequency. This allows the system to respond just in time adapting to the existing traffic conditions rather than conservatively allocating more guard channels. Thus, it is possible to use both reassignments and guard channels synergistically. Reassignments mitigate the negative effect of guard channels on bandwidth utilization while guard channels prevent excessive reassignments.

We shall now see two such adaptive schemes LOREL and MXREV that utilize the knowledge of reassignments in the neighborhood differently. LOREL attempts to contain the number of reassignments under a specified limit while MXREV tries to maximize the revenue by balancing the penalty for reassignments with the reward for serviced calls. From a quality of service point of view it is desirable to place an upper bound on the number of times a mobile is forced to switch channels. From a service provider's point of view it is desirable to maximize the revenue generated by the service. LOREL and MXREV respectively address these two perspectives.

In both these schemes a cell uses its local neighborhood information in deciding the number of guard channels. It is assumed that each cell is aware of the number of reassignments, the number of admitted calls and the number of blocked calls in the neighborhood. Based on this information each cell periodically adjusts its number of guard channels. It is assumed that this period is big enough to capture steady state behavior but much smaller than the time between changes in traffic conditions. Note that these schemes are characterized by distributed decision making using local information. While each cell acts independently, collectively they have the effect of achieving a common objective. Following subsections describe these schemes in detail.

A. BOUNDING REASSIGNMENTS

The bounded reassignment scheme LOREL attempts to contain reassignment frequency below a specified limit r^* as follows. Each cell periodically adjusts the number of guard channels g^*, based on reassignment frequency in the neighborhood. It keeps track of the number of calls admitted, c, and the number of reassignments, r, in the neighborhood since the last adjustment to its g^*. If either the number of new call requests in that cell reaches a threshold value n' or r reaches a threshold value r', then g^* is recomputed as follows.

$$g^* \leftarrow \begin{cases} g^* + 1, & r/c > r^*, \\ \max(g^* - 1, 0), & r/c < r^* - \delta, \\ g^*, & otherwise. \end{cases} \tag{1}$$

In each period the number of guard channels is increased if the reassignment frequency is above the given limit r^* and it is decreased if the frequency is below $r^* - \delta$. When the reassignment frequency is within the range $(r^* - \delta, r^*)$ the number of guard channels remains unchanged. The parameter δ is used to insure against under-utilization by not allowing the number of reassignments to fall too far below the specified limit. By making cells recompute their g^* values whenever the number of reassignments reaches the threshold value, LOREL ensures that corrective action is taken whenever congestion is seen in the neighborhood. On the other hand, by recomputing g^* periodically based on the number of new call arrivals, it ensures that the bandwidth does not go underutilized because of excess guard channels.

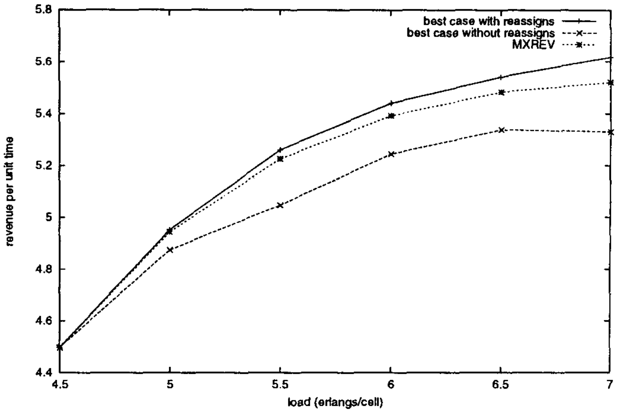

Fig. 6. Revenue as a function of load

Figure 5 compares the performance of LOREL with GUARD and HOREL. For the purpose of this experiment, we chose, $r^* = 0.02$, $\delta = 0.005$, $n' = 25$, and $r' = 20$. Thus, the number of guard channels g^* in a cell is adjusted after every 25 new call arrivals in that cell or after 20 reassignments in its neighborhood. If the reassignment frequency is above 0.02, g^* is incremented and it is decremented if reassignment frequency falls below 0.015. g^* stays unchanged if the reassignment frequency is between 0.015 and 0.02. From the figure, it can be seen that reassignment frequency is contained under the given limit of 0.02 at all loads. Comparing the blocking probabilities and reassignment frequencies of LOREL with that of HOREL, it can be said that LOREL achieves substantial reduction in the reassignment frequency for a relatively small increase in the blocking probability. Compared to GUARD, the blocking probability of LOREL is significantly lower at low loads and the difference narrows at high loads. Further, it should be noted that there were no dropped calls under LOREL. To sum up, the combination of controlled reassignments and dynamically determined guard channels performs well by adapting to the traffic conditions.

B. MAXIMIZING REVENUE

Reassignments effectively eliminate handoff failures and also make it possible for the system to be less conservative in reserving guard channels resulting in higher utilization. But reassignments may inconvenience users and incur communication and processing overheads. A service provider would like to use reassignments only if they are profitable. From this point of view, an ideal scheme is one that maximizes the revenue V, given that each call serviced brings in reward R and each reassignment incurs penalty P. We now describe a scheme, called MXREV, that attempts to maximize the revenue given R and P. We do not consider the penalty for call dropping since reassignments almost completely eliminate them.

Each cell keeps track of the number of new calls blocked b and the number of reassignments r in its neighborhood. To maximize the revenue, MXREV tries to minimize L, the sum of the penalty due to reassignments and the loss in revenue due to blocked calls, $b \times R + r \times P$, by dynamically adjusting the number of guard channels, g^*. Note that changing g^* has opposing effects on b and r. As in the case of LOREL, the guard channels are adjusted whenever the r reaches a threshold r' or the number of new call

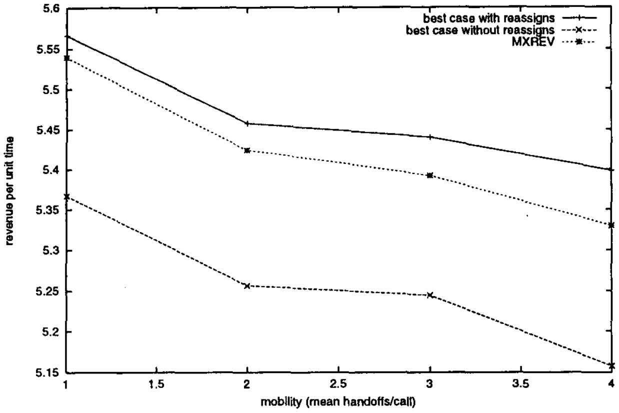

Fig. 7. Revenue as a function of mobility

requests n reaches a threshold n'.

$$g^*_{i+1} \leftarrow \begin{cases} 2 \times g^*_i - g^*_{i-1}, & L_i < L_{i-1}, \\ g^*_{i-1}, & otherwise. \end{cases} \tag{2}$$

The cells remember the values of L and g^* for the last two periods i and $i-1$. If the loss in revenue decreased from L_{i-1} to L_i, the number of guard channels, g^* will be moved in the same direction for the next period $i+1$ as was done in transitioning from period $i-1$ to i. Otherwise, g^* will be moved in the opposite direction. In any case, g^* is always moved in steps of 1.

Intuitively, this scheme performs a blind hill-climbing. For a given load and mobility, the revenue peaks at a specific number of guard channels. But due to changing traffic conditions this is a moving peak. In general, it is not possible to determine if the peak has been reached. The scheme reverses direction whenever it senses that it is moving down the hill (when loss in revenue increases) and moves in the same direction as long as it is climbing the hill (when loss in revenue decreases). Thus the number of guard channels hovers around the optimal value for the current traffic conditions.

Figure 6 compares the revenue generated by MXREV with the best observed revenues with and without reassignments under varying load when both P and R are 1. The best revenue with reassignments was obtained for each load by selecting the maximum revenue from the experiments with different guard channels. The best revenue without reassignments was computed by considering the case that achieved a dropping probability less than 1/10000 with the minimum number of guard channels. This limit on dropping probability was chosen for fair comparison since the dropping probability for MXREV is almost zero. Figure 7 shows a similar comparison under varying mobility. In both these figures, the difference in revenue between the two best cases shows the potential benefit from using reassignments. The figures show that MXREV closely approximates the performance of the best case with reassignments. Fixed guard channel schemes and even adaptive schemes that do not use reassignments are unlikely to achieve the performance of the best case without reassignments (the bottom curve). The use of reassignments allows the luxury of reacting to the traffic conditions just in time. This explains the superior performance of MXREV.

Note that the revenue increases with increasing load and gradually flattens out. As the offered load increases the amount of calls serviced also increases. But the capacity to admit more calls decreases gradually and hence the incremental gains tend to grow smaller. On the contrary, the revenue decreases with increasing mobility. Due to increase in the number of handoffs more guard channels are required thus increasing blocking probability. The number of reassignments also increases. Each contributes to the decline in revenue.

V. CONCLUDING REMARKS

We showed that reassignments when used in a controlled manner increase the utilization without sacrificing the quality of service. We also showed how guard channels can be used to control reassignments. We presented two schemes that are based on reassignments which dynamically adjust guard channels in each cell adapting to the local traffic conditions. Simulation results showed that these schemes perform well under uniform traffic conditions. The behavior of these schemes under non-uniform traffic conditions with hot spots is currently under investigation. Issues involved in employing these schemes in real cellular systems need to be addressed. An important aspect of these schemes is their local decision making in fixing the number of guard channels. Its effect on fairness in resource distribution among different cells must be analyzed.

REFERENCES

[1] D. Cox, and D.Reudnik, "Increasing Channel Occupancy in Large-Scale Mobile Radio Systems: Dynamic Channel Reassignment", *IEEE Transactions on Communications*, 21(11), November 1973.

[2] Chih-Lin I and Pi-Hui Chao, "Local Packing - Distributed Dynamic Channel Allocation at Cellular Base Station," *IEEE GLOBECOM '93*.

[3] D. Hong and S.S. Rappaport, "Traffic Model and Performance Analysis for Cellular Mobile Radio Telephone Systems with Prioritized and Nonprioritized Handoff Procedures," *IEEE Transactions on Vehicular Technology*, 35(3), 1986.

[4] I. Katzela and M. Naghshineh, "Channel Assignment Schemes for Cellular Mobile Telecommunication Systems: A Comprehensive Survey", *IEEE Personal Communications*, June 1996.

[5] D. Levine, I. Akylidz, and M. Naghshineh, "A Resource Estimation and Call Admission Algorithm for Wireless Multimedia Networks Using the Shadow Cluster Concept," *IEEE/ACM Transactions on Networking*, 5(1), February 1997.

[6] M. Naghshineh and M. Schwartz, "Distributed Call Admission Control in Mobile/Wireless Networks," *IEEE Journal on Selected Areas in Communications*, 14(4), May 1996.

[7] S. Nelakuditi, Z. Zhang, and D. Du, "Reassignment based Call Admission Control for Cellular Networks," *IEEE LANMAN'98*, May 1998.

[8] R. Ramjee, D. Towsley, and R. Nagarajan, "On Optimal Call Admission Control in Cellular Networks," *ACM/Baltzer Journal of Wireless Networks*, 3(1), 1997.