Hot-Spot Traffic Relief in Cellular Systems

Tak-Shing Peter Yum and Wing-Shing Wong

Abstract— **During the operation of a cellular system, unexpected growth of traffic may develop in various cells and create local traffic congestions. Channel borrowing, at least in a "semi-dynamic" way, is known to be a viable solution to alleviate congestions. In this paper, we show by analysis of mathematical models that combining channel borrowing with a coordinated sectoring or overlying scheme provides effective ways to handle hot-spots in the system. Blocking probabilities with these arrangements are derived, and the Dynamic Sharing with Bias (DSB) rule is suggested for increasing the trunking efficiency. A simple handoff model is formulated and analyzed for comparing the probabilities of additional handoffs due to sectoring and overlaying of cells. With the nominal allocation of 60 channels per cell and a donor cell having a load of 30 Erlangs, numerical results show that at a blocking requirement of 1 percent, the traffic load in the hot-spot cell can be increased from 47 to 63 Erlangs with the use of the channel borrowing with the cell sectoring scheme; while with the use of the DSB rule, the load can be increased further to 71 Erlangs. A slightly higher load can be carried in the hot-spot cell with the use of cell overlaying arrangement.**

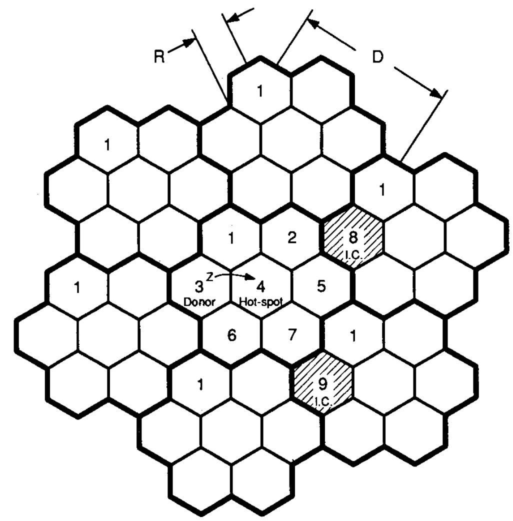

Fig. 1. The lending of channel set Z from the donor cell to the hot-spot cell and the locking of Z on the interfering cells (I.C.).

I. Introduction

The engineering of a mobile cellular telecommunication system is a complex process requiring huge quantities of measured and forecast data, and involving lengthy and complicated computations [1]. Even under the most ideal situations, as a mobile cellular system grows uneven traffic load may develop into "hot-spot" cells, that is, cells with traffic load substantially larger than the design load. To alleviate this kind of unexpected load on the hot-spot cells, several approaches can be taken, including: cell splitting, channel borrowing, cell sectoring, and cell overlaying.

Cell splitting [2], [3] is the most effective solution when there is a large number of hot-spot cells. However, for a system with a few isolated hot-spot cells, this approach is very expensive since additional equipment, such as radio transceivers, antennas, power plant, data terminals, and MTSO interfaces, are needed for each new cell. Moreover, many cells need to be split simultaneously before additional channel resources are created.

On the other hand, channel borrowing from adjacent cells is less costly. It is an effective approach when the traffic loads at the neighboring cells are light. The penalty of channel borrowing, however, is very large in that to avoid interference, the lending cell as well as its two co-channel cells closest to the hot-spot cell are forbidden to use the channels being lent [4]. We illustrate this case in Fig. 1. Also, if the traffic loads at the neighboring cells are not light, indiscriminate borrowing of channels will deplete the channel resources and induce congestions elsewhere.

Sectoring is used in some existing commercial systems. It is used to increase channel reuse and is applied to a group of neighboring cells. Since new equipment is needed, it is an expensive idea to upgrade an omnicell system to a sectorized system. On the other hand, our approach in this paper of using borrowing in sectorized cells only requires sectoring the hot-spot cell and the two interfering cells. This is economically more desirable, especially if one is not sure whether the existing hot-spot is a permanent or a temporary phenomenon.

In this paper, we study the effect of relieving congestion at a single hot-spot cell by using cell sectoring [5] and cell overlaying [3], [6], [7]. We examine the blocking probability of fixed channel allocation methods associated with these schemes and propose a very efficient channel sharing method.

Although we assume the hot-spots are isolated in the paper, our method clearly can be extended to the less likely scenario where there are multiple hot-spots. To handle this case, we can assume the hot-spots do not occur simultaneously, but in a sequential order. Our algorithm can then be applied recursively to this scenario.

The sectoring and the overlaying of cells induce more call handoffs (also called handover in some literature [8], [9]). In [10], a simple handoff activity model was proposed and

Manuscript received March 16, 1992; revised November 19, 1992. This research was done while the authors were at the Performance Analysis Department of AT&T Bell Laboratories. This paper was presented at the IEEE ICC '92, Chicago, IL, June 1992.
The authors are with the Department of Information Engineering, The Chinese University of Hong Kong, Shatin, Hong Kong.
IEEE Log Number 9208655.

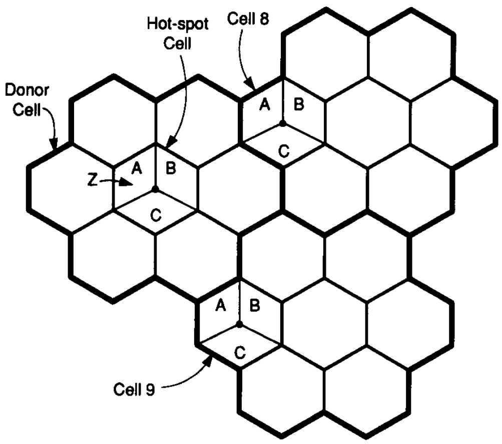

Fig. 2. Interference prevention through cell sectoring.

the probability of additional handoff due to cell sectoring cell overlaying was derived. It was shown that the more efficient channel allocation method has a larger probability of needing an additional handoff. Thus, to reengineer the hot-spot cells properly, one has to consider the impact to the processing requirements due to additional handoffs.

II. Channel Borrowing in Sectorized Cells

For the comparison of trunking efficiency of various borrowing schemes, consider a cellular system with uniform size cells of radius R and a regular seven cell channel reuse pattern as shown in Fig. 1.

Without loss of generality, let cell 4 be a hot-spot cell in the system and let the traffic load in cell 3, a neighbor of cell 4, be relatively light so that the lending of some channels to cell 4 is possible without affecting its grade of service. We shall, for convenience, call cell 3 the donor cell. Let X be the set of channels nominally allocated to the hot-spot cell, let Z be the set of channels relocated from the donor cell to the hot-spot cell, and let Y be the remaining set of channels in the donor cell. For an arbitrary channel set S, let $|S|$ denote the total number of channels in S.

Cells 8 and 9 are the co-channel cells of the donor cell and can normally use channel sets Y and Z. We shall, for convenience, refer to them as the interfering cells. There will be unacceptable interference if both the hot-spot cell and the interfering cells use channel set Z simultaneously. To prevent the hot-spot cell from interfering with cells 8 and 9, we sector the hot-spot cell as shown in Fig. 2 and let channel set Z be used only in sector A. Similarly, to prevent cells 8 and 9 from interfering with the hot-spot cell, we also sector these interfering cells and allow only channel set Y be used in sector A.

As the hot-spot cell can use both channel set X and Z, it is possible to design various algorithms to allocate channels to calls. We assume the following allocation scheme is used.

1) *Calls in Sector A:* Use channels from set Z first. If none is available, use channels from set X.

2) *Calls in Sectors B and C:* Use channels from set X only.

This scheme allows the maximum number of channels available for calls in sectors B and C at the expense of seeing a higher possibility of additional handoffs.

In the rest of this section, we will derive the blocking probabilities of the system under such an allocation scheme. To make the analysis simple, we assume that the call arrivals are Poisson with rates $\lambda_D, \lambda_I,$ and λ_H in the donor cell, the interfering cell, and the hot-spot cells, respectively. We also assume that the call arrival density is uniform in a cell, and that the sectors are symmetrically divided. Hence, the traffic rates are identical among the sectors in a cell. Finally, we assume the call durations to be exponentially distributed with mean $1/\mu$.

A. Blocking Probability at the Interfering Cells

Before cell sectoring, the blocking probability B_0 is given by the Erlang B formula

$$B_0 = E_2(\lambda_I, |Y| + |Z|)$$

where

$$E_2(\lambda, N) = \frac{\frac{(\lambda/\mu)^N}{N!}}{\sum_{n=0}^{N} \frac{(\lambda/\mu)^n}{n!}}.$$

After sectoring, calls in sector A are assigned only to channels in Y, while calls in sectors B or C are assigned to channels in Z first, and then to channels in Y. Since there are restrictions on the use of certain channels in certain sectors, the call blocking probability will be somewhat increased.

Let $S_{j,k}$ denote the state where j channels in Y and k channels in Z are occupied in the interfering cell. According to the above channel allocation scheme, the state transition diagram as shown in Fig. 3 can be obtained with $p = 1/3$. The state probabilities can be solved numerically using conventional techniques. The call blocking probability in sector A, denoted as $B_I(A)$, is just the probability that all channels in set Y are used up and is given by

$$B_I(A) = \sum_{k=0}^{|Z|} P[S_{|Y|,k}].$$

Calls initiated in sectors B or C are blocked when channels in sets Y and Z are all used up. Denoting this probability as $B_I(B \cup C)$, we have

$$B_I(B \cup C) = P[S_{|Y|,|Z|}].$$

The overall blocking probability in the interfering cell, given that the traffic rates are the same in each sector, is

$$B_I = \frac{1}{3} B_I(A) + \frac{2}{3} B_I(B \cup C).$$

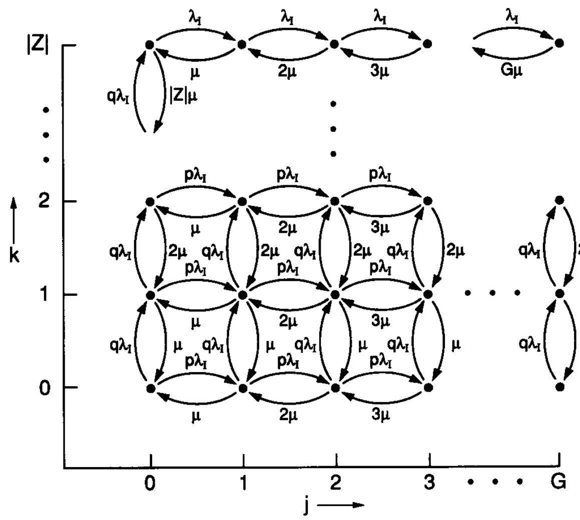

Fig. 3. State transition rate diagram, $q = 1 - p$.

B. Blocking Probability at the Hot-Spot Cell

Let $S'_{i,k}$ be the state that i channels in X and k channels in Z are occupied. According to the above assignment strategy, we can obtain the state transition diagram shown in Fig. 3 with j replaced by i and $p = 2/3$. The blocking probabilities are similarly obtained as

$$B_H(A) = P\left[S'_{|X|,|Z|}\right]$$

$$B_H(B \cup C) = \sum_{k=0}^{|Z|} P\left[S'_{|X|,k}\right]$$

$$B_H = \frac{1}{3} B_H(A) + \frac{2}{3} B_H(B \cup C).$$

The call blocking probability in the donor cell is simply

$$B_D = E_2(\lambda_D, |Y|).$$

III. Dynamic Channel Sharing

Very often, the donor and hot-spot cells are served by the same MTSO. In that case, dynamic sharing of channel set Z among these two cells can be implemented easily. In this section, we investigate the potential improvement of trunking efficiency with this kind of sharing.

Various ways of sharing are possible [8], [10]. The optimal sharing strategy appears to be very difficult to obtain since in general it should depend on the channel occupancies in X, Y, and Z, the channel set sizes, as well as on the traffic load λ_D and λ_H. What we propose instead is a rule called Dynamic Sharing with Bias (DSB). The DSB rule is similar to the Join-Biased-Queue rule [11], whereby channel assignment decision depends only on the number of the available channels from the channel sets, which can be tracked easily. Queueing analysis shows that the Join-Biased-Queue rule is simple but effective in balancing the load of servers in the presence of unbalanced traffic streams. Numerical results in Section VI show that the DSB rule with an optimized bias parameter Δ is equally effective. Specifically, the rule says the following.

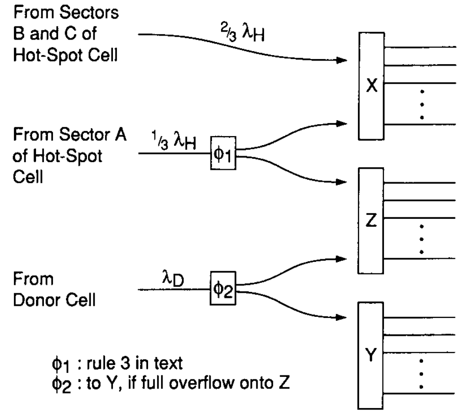

Fig. 4. Dynamic sharing of channel set Z.

1) *Calls in Donor Cell:* Assign channels in Y first, then assign channels in Z.

2) *Calls in Sectors B and C of the Hot-Spot Cell:* Assign channels in X.

3) *Calls in Sector A of the Hot-Spot Cell:* Assign channels in X or Z so that the remaining number of channels available in the hot-spot cell is as close as possible to the remaining number of channels available in the donor cell plus Δ. More specifically, if i, j, and k represent the number of channels already used in the channel sets X, Y, and Z, respectively, then

a) if ($i = |X|$ and $k < |Z|$) or ($i < |X|$ and $k = |Z|$), assign a channel in the nonblocking channel set, and

b) if ($i < |X|$ and $k < |Z|$), assign a channel in X or Z according to the following decision rule:

$$(|X| - i) \overset{X}{\underset{Y}{\gtrless}} (|Y| + |Z| - j - k) + \Delta.$$

Rule 3(b) says that if the remaining number of channels in X is larger than Δ plus the combined number of free channels in Y and Z, assign a channel in X; if smaller, assign a channel in Z. Δ can be chosen as a noninteger to resolve the equality case. Fig. 4. shows graphically the DSB rule.

To minimize blocking, the optimal value of Δ should obviously be a function of the state variables i, j, and k as well as λ_D and λ_H. We observe that when the remaining numbers of available channels in the two cells are large, the choice of Δ is not critical as long as Δ is not too large. When the remaining numbers of available channels are small, then a small value of Δ is sufficient for biasing. Therefore, for minimum blocking, the optimal choice of Δ should not be very sensitive to the state variables. What we will do is to assume that it is independent of the state variables and search for the optimal Δ for a give λ_D and λ_H. In practice, optimal Δ's can be computed off-line and stored. When a new estimate of λ_D and λ_H is obtained, the optimal Δ is lookup and used.

Let $S''_{i,j,k}$ denote the state where i channels in X, j channels in Y, and k channels in Z are occupied. As in the previous

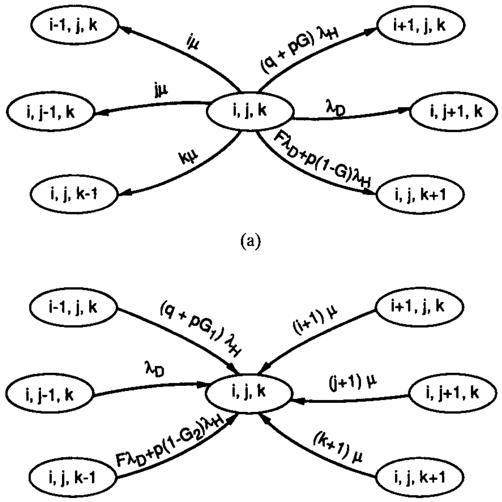

Fig. 5. Transition rates out (a) and in (b) of state $(i, j, k), p = 1/3$, $q = 1 - p, F, G, G_1$, and G_2 are binary control variables.

cases, a Markov chain for the state transition process can be formed and solved numerically. As the chain is now three-dimensional, the channel set sizes cannot be too large for practical solutions. Fig. 5 shows the transition rates in and out of a typical state $S''_{i,j,k}$. The call blocking probability is just the sum of the probabilities of the blocking states. Specifically,

$$B_D = \sum_{i-0}^{|X|} P\left[S''_{i,|Y|,|Z|}\right]$$

$$B_H(A) = \sum_{j=0}^{|Y|} P\left[S''_{|X|,j,|Z|}\right]$$

$$B_H(B \cup C) = \sum_{k=0}^{|Z|} \sum_{j=0}^{|Y|} P\left[S''_{|X|,j,k}\right].$$

IV. Channel Borrowing with Cell Overlaying

Besides sectoring, another way to reduce interference is to underlay a cell of smaller size [3], [6], [7]. Channels used in the underlaid cell can therefore be reused in a smaller distance. To compare cell overlaying with cell sectoring, let us consider the same scenario as before, where a regular seven cell channel reuse pattern is defined. We assume that cell 4 is the hot-spot cell and channel set Z is borrowed from cell 3. Let this borrowed set of channels be used in the underlaid hot-spot cell. To eliminate the mutual interference between the underlaid hot-spot cell and interfering cells (Fig. 6), the interfering cells must also be partitioned into the overlay–underlay arrangement where channel set Z can only be used in the underlaid cells. The call blocking analysis in the hot-spot cell is exactly the same as that for cell sectoring arrangement with the underlaid cell corresponding to sector A and the overlaid cell corresponding to sectors B and C. For the interfering cells, the correspondence is reversed, i.e., the underlaid cell corresponds

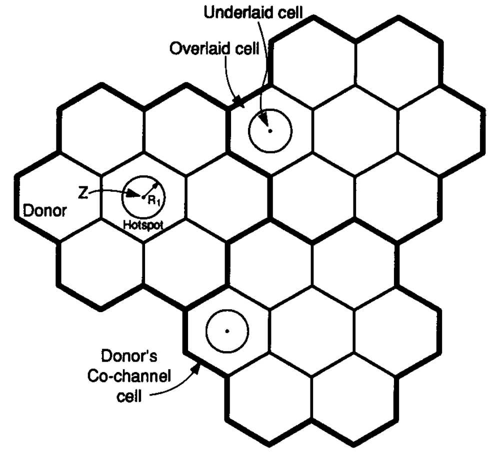

Fig. 6. Interference prevention through cell overlaying.

to sectors B and C and the overlaid cell corresponds to sector A, for obvious reason. Because we have assumed uniform traffic density within a cell, the probability that a call is initiated in the underlaid cell of radius R_1 denoted as f, is given as

$$f = \frac{\text{area of an underlaid cell}}{\text{area of a whole cell}} = \frac{\pi R_1^2}{(3\sqrt{3}/2)R^2}.$$

Therefore, the traffic load at the underlaid and overlaid cells are, respectively $f\lambda$ and $(1 - f)\lambda$ for a cell with traffic load λ. The derivation of blocking probabilities with and without channel sharing is exactly the same as that for a sectored cell.

A key parameter in determining the reuse distance is defined as

$$q \equiv \frac{\text{shortest distance between two co-channel cells}}{\text{cell radius}}.$$

For the seven cell reuse pattern assumed here, q can be computed, using the cosine law of trigonometry, to be

$$q = \sqrt{12 + 3 - 12\cos\left(\frac{2\pi}{3}\right)} = 4.58.$$

To preserve $q = 4.58$ between the underlaid hot-spot cell and the underlaid interfering cell, R_1 must not be too large. The distance between the hot-spot cell and cell 8 is $3R$, and that between the hot-spot cell and cell 9 is $2\sqrt{3}$ or $3.46R$. Therefore, R_1 must be selected such that

$$q = \frac{3R}{R_1} \geq 4.58$$

or

$$R_1 \leq 0.655R.$$

Hence, to preserve $q \geq 4.58$, the maximum fraction of area F that can be covered by the underlaid cell is

$$F = \frac{\pi(0.655R)^2}{(3\sqrt{3}/2)R^2} = 0.519.$$

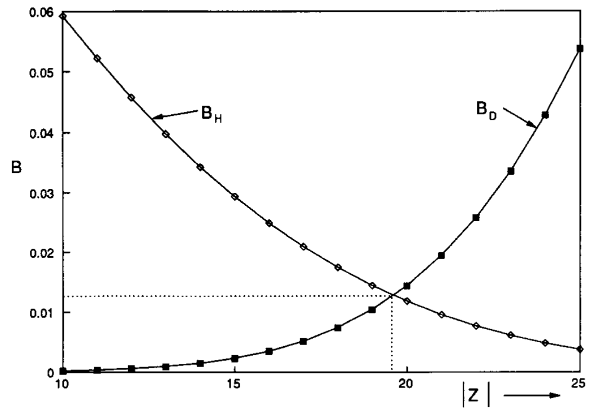

Fig. 7. Channel borrowing in sectorized cells λ_H = 66 Erlang, λ_D = 30 Erlang,

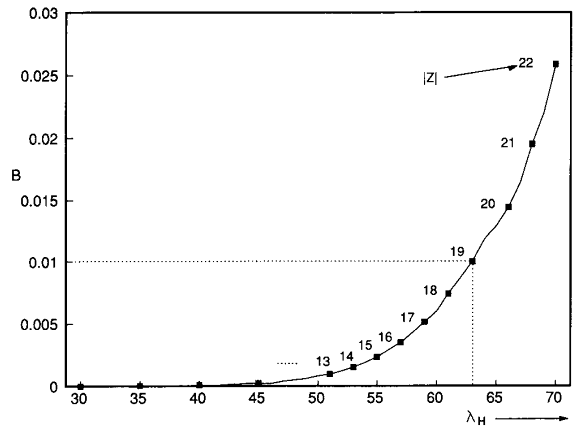

Fig. 8. Channel borrowing in sectorized cells, $|X| = 60$, $\lambda_D = 30$ Erlang, $B = \max(B_H, B_D), B_H \approx B_D$.

V. Numerical Results

In the following numerical examples, we assume μ is normalized to unity so that the traffic rate λ can be expressed in Erlangs. Consider a cellular system where each cell has a nominal number of 60 channels. The nominal peak load at a call blocking probability of 10^{-2} is 47 Erlangs from the Erlang B formula. Let us assume the hot-spot cell has a traffic load of 66 Erlangs and the donor cell a load of 30 Erlangs. Let us assume that $|Z|$ channels are reassigned from the donor cell to sector A of the hot-spot cell. The resulting blocking probabilities B_H and B_D in the two cells as a function of $|Z|$ are shown in Fig. 7. It is seen that the optimal choice of $|Z|$ is 20 at $B_H \simeq B_D \simeq 0.014$. In fact, there is an optimal $|Z|$ for every value of λ_H under the condition of $B_H \approx B_D$. Fig. 8 shows $B \equiv \max(B_H, B_D)$ versus λ_H using the optimal $|Z|$. It is seen that at a blocking requirement of $\max(B_H, B_D) \leq 10^{-2}$, λ_H can be as high as 63 Erlangs at $|Z| = 19$. The blocking probability in sector A is 0.42 percent, while that in sectors B and C is 1.07 percent. The average blocking probability B_H is 0.85 percent. Compared to a carrying capacity of 64.4 Erlangs for a cell with 79 (69 + 19) channels but without sectoring (i.e., the pure channel borrowing case), this represents a very small loss of trunking efficiency. The benefit of sectoring is that all co-channel cells can reuse the borrowed set of channels. To say it in another way, the above case shows that when a neighbor of the hot-spot cell has a load which is 36 percent [(47–30)/30] under capacity, channels can be borrowed from it to alleviate congestions in the hot-spot cell up to a load of 34 percent [(63–47)/47] above the designed load.

To see the effect of sectoring on the interfering cells, consider cell 8 in Fig. 2. Here, calls in sector A are assigned only to channels in set Y while calls in sectors B and C are assigned to channels in Z first and then to channels in Y. Fig. 9 compares the blocking probabilities of an interfering cell before and after sectoring assuming the traffic distribution is uniform on the entire cell and $|Z| = 12$. As an example, at a load of 47 Erlangs, the blocking probability is 0.01

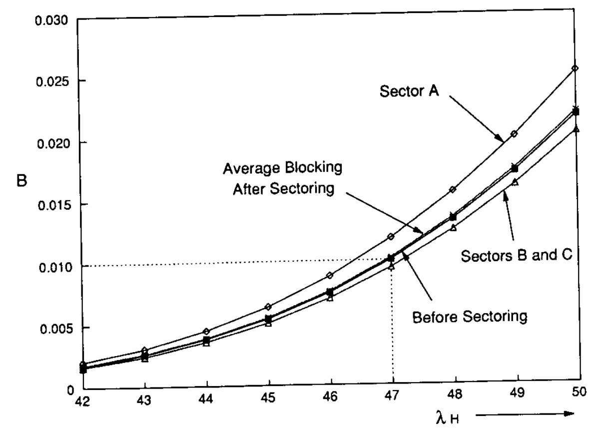

Fig. 9. Blocking comparison in the interfering cells with $|Z| = 12$.

before sectoring in the interfering cell. Fig. 9 shows that after sectoring, the blocking rate in sector A is slightly increased to 1.2 percent, while that in sectors B and C is slightly decreased to 0.9 percent. The average, however, is still 1 percent (0.33 × 1.2 percent + 0.67 × 0.9 percent). Similar performance is obtained when $|Z|$ is increased to 19. The effect of sectoring on the interfering cell is therefore minimal.

To investigate the potential benefit of dynamic channel sharing, we consider the similar case (i.e., $\lambda_H = 50$, $\lambda_D = 30$) and choose Δ to minimize

$$L = \lambda_D B_D + \lambda_H \left[\frac{1}{3} B_H(A) + \frac{2}{3} B_H(B \cup C) \right]$$

where L is the expected number of blocked calls per second in the two cells. Fig. 10 shows L versus Δ. At the optimal value of $\Delta = 1$, B_D and B_H are computed to be 5.1×10^{-5} and 7.3×10^{-5}, respectively, and are both substantially less than that without sharing. Actually, λ_H can be increased to 71 Erlangs at $\Delta = 5$ without violating the $B \leq 10^{-2}$ require-

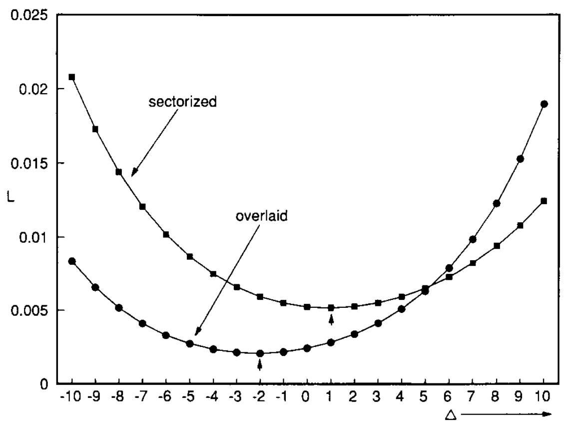

Fig. 10. Optimal Δ in sectorized and overlaid arrangements with dynamic channel sharing, $|X| = 60$, $|Z| = 30$, $\lambda_H = 50$ Erlang, $F = 0.519$.

ment. Compared to the nominal peak load of 47 Erlangs, this represents a 51 percent increase of system capacity.

As mentioned before, for congestion relief through cell overlaying, the results are exactly the same as that for cell sectoring with $F = 0.333$. For designs with the underlaid cell at maximum size, or at $F = 0.519$, the call carrying capacity of the hot spot-cell can be increased to 64 Erlangs in the previously defined scenario. This 1.6 percent increase over the cell sectoring arrangement is due to the fact that more traffic can use the borrowed channels in the cell overlaid arrangement. Using dynamic channel sharing, Fig. 10 shows that Δ is at -2 with $B_D = 3.7 \times 10^{-5}$ and $B_H = 1.9 \times 10^{-5}$. At $B \leq 10^{-2}$, λ_H can be increased to 75 Erlangs at $\Delta = 0$. Thus, for this example, a 60 percent increase of nominal load can be tolerated with the Dynamic Channel Sharing.

VI. Conclusions

We have shown that in a cellular system, spare channels in neighboring cells can be borrowed to relieve local congestions. These borrowed channels should be used in such a way that interference in co-channel cells is avoided. Cell sectoring and cell overlaying are known to be effective ways to reduce interference among co-channel cells. We found that such arrangements can substantially increase the capacity of the congested cells. The capacity can be further increased with the use of Dynamic Channel Sharing. A model for comparing handoff probabilities with different channel assignments was proposed in [10]. It was found that the cell sectoring arrangement has a smaller probability of requiring an additional handoff. The tradeoff, however, is that the sectorized arrangement requires additional directional antennas, while the same omnidirectional antenna can be shared by the two concentric cells in the overlaid arrangement.

The above cell and channel engineering methodology is equally applicable to digital cellular system where the unit of borrowing is a multichannel carrier. Also, we have demonstrated the borrowing of channels from one neighbor. In practice, as many as six neighbors can contribute their resources for the relief of a local congestion. We have not examined the problem of multiple hot-spots here. We believe the next challenging problem is to design some algorithms to optimally borrow channels from multiple donors and to relieve multiple hot-spots.

Acknowledgment

It is a pleasure to acknowledge the many helpful discussions we have had with S. Nanda, J. Kauffman, R. Yeung, and O.-C. Yue. Our sincere thanks are also due to M. Chen and L. Yeung who are responsible for all the numerical computations.

References

[1] W. C. Y. Lee, *Mobile Cellular Telecommunications Systems.* New York: McGraw-Hill, 1989.
[2] R. H. Frenkiel, "A high-capacity mobile radio telephone system model using a coordinated small-zone approach,"*IEEE Trans. Veh. Technol.,* vol. VT-19, May 1970.
[3] J. F. Whitehead, "Cellular spectrum efficiency via reuse planning," in *IEEE Veh. Technol. Conf. Rec.,* 1985, pp. 16–20.
[4] M. Zhang and T. S. Yum, "Comparisons of channel assignment strategies in cellular mobile systems," *IEEE Trans. Veh. Technol.,* Nov. 1989.
[5] V. H. MacDonald, "The cellular concept," *Bell Syst. Tech. J.,* vol. 58, no. 1, Jan. 1979.
[6] S. W. Halpern, "Reuse partitioning in cellular systems," in *IEEE Veh. Technol. Conf. Rec.,* May 1983.
[7] W. C. Y. Lee, "New cellular schemes for spectral efficiency," *IEEE Trans. Veh. Technol.,* vol. VT-36, Nov. 1987.
[8] R. Beck and H. Panzer, "Strategies for handover and dynamic channel allocation in microcellular mobile radio," in *IEEE Veh. Technol. Conf. Rec.,* 1989, pp. 668–672.
[9] D. Hong and S. S. Rappaport, "Traffic model and performance analysis for cellular mobile radio telephone systems with prioritized and nonprioritized handoff procedures," *IEEE Trans. Veh. Technol.,* vol.-35, Aug. 1986.
[10] K. Sallberg, B. Stavenow, and B. Eklundh, "Hybrid channel assignment and partitioning in a cellular mobile telephone system," in *IEEE Veh. Technol. Conf. Rec.,* June 1987, pp. 405–411.
[11] T. S. Yum and M. Schwartz, "The join-biased-queue rule and its application to routing in computer communication networks," *IEEE Trans. Commun.,* Apr. 1981.
[12] T. S. Yum and W. S. Wong, "Hot-spot traffic relief in cellular systems," in *Proc. IEEE Int. Conf. Commun.,* Chicago, IL, June 1992, pp. 1703–1709.

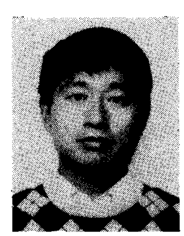

Tak-Shing Peter Yum received the B.Sc., M.Sc., M.Ph., and Ph.D. degrees from Columbia University, in 1974, 1975, 1977, and 1978, respectively.

He worked at Bell Telephone Laboratories for $2\frac{1}{2}$ years and taught in National Chiao Tung University, Taiwan, for 2 years before joining The Chinese University of Hong Kong in 1982. He has published original research on packet switched networks with contributions in routing algorithms, buffer management, deadlock detection algorithms, message resequencing analysis, and multiaccess protocols. In recent years, he branched out to work on the design and analysis of cellular network, lightwave networks, and video distribution networks. He believes that the next challenge is designing an intelligent network that can accommodate the needs of individual customers.

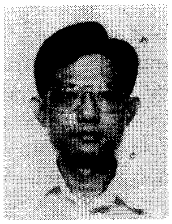

Wing Shing Wong received the combined M.A.B.A. degree from Yale University, New Haven, CT, in 1976. He received the M.S. degree in 1978 and the Ph.D. degree in 1980, both from Harvard University, Cambridge, MA.

He joined AT&T Bell Laboratories in 1982 and was promoted to Supervisor in 1987. He is currently a Reader in Information Engineering of the Chinese University of Hong Kong. He has served as a consultant on performance and design projects ranging from UNIX systems to cellular telecommunications systems. His research interests include performance evaluation, neural computing, combinatorial optimization, and nonlinear filtering.