Title
Worst-case performance of cellular channel assignment policies

Journal
Wireless Networks, 2(4)

Authors
Jordan, Scott
Schwabe, Eric

Publication Date
1996-12-01

DOI
10.1007/BF01262046

Peer reviewed

Baltzer Journals February 21, 1996

Worst-Case Performance of Cellular Channel Assignment Policies

Scott Jordan* and Eric J. Schwabe†
Department of EECS
Northwestern University
2145 Sheridan Road
Evanston, IL 60208
E-mail: {scott,schwabe}@eecs.nwu.edu

Many cellular channel assignment policies have been proposed to improve efficiency beyond that resulting from fixed channel allocation. The performance of these policies, however, has rarely been compared due to a lack of formal metrics, particularly under nonhomogeneous call distributions. In this paper, we introduce two such metrics: the worst-case number of channels required to accommodate all possible configurations of N calls in a cell cluster, and the set of cell states that can be accommodated with M channels. We first measure two extreme policies, fixed channel allocation and maximum packing, under these metrics. We then prove a new lower bound, under the first metric, on any channel assignment policy. Next, we introduce three intermediate channel assignment policies, based on commonly used ideas of channel ordering, hybrid assignment, and partitioning. Finally, these policies are used to demonstrate the tradeoff between the performance and the complexity of a channel allocation policy.

1 Introduction

Wireless services are one of the strongest growth areas in telecommunications today. Cellular voice is well established as a high-end service in most areas, but demand is increasing rapidly. Personal communications services (PCS) are expected to be introduced in the next few years as a mass market phone service. Wireless data services are appearing in the form of cellular digital packet data (CDPD), wireless local area networks (LANs), and wireless modems. Capacity, however, is now a critical issue for all of these services. In response, carriers are investigating cell splitting, allocation of new spectrum, alternative multiple access architectures, and dynamic channel allocation. In this paper, we focus on this last approach for increasing the available capacity.

*Supported by the Ameritech Foundation.
†Supported by the National Science Foundation under grant CCR–9309111.

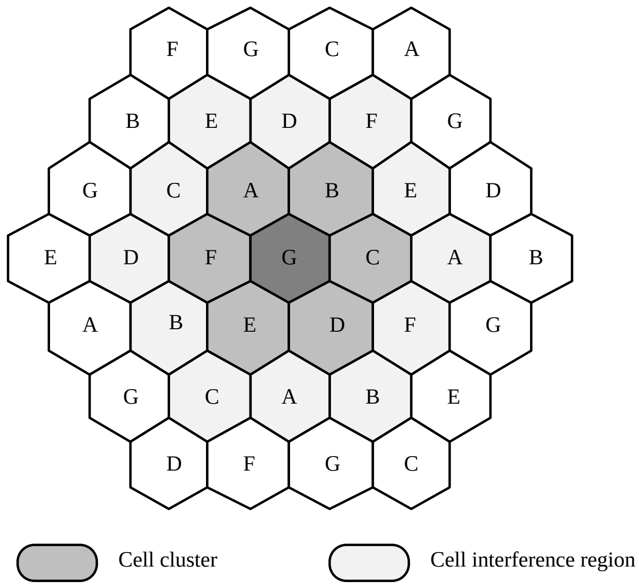

Figure 1: A cell cluster and interference region.

In cellular systems, the geographical region is split, using a regular topology, into *cells*, each containing one base station. The most common cell shape in two dimensions is a regular hexagon (Figure 1). A mobile wishing to initiate a call must request a channel from the base station in the cell in which the mobile currently exists. The base station must assign a channel that is not currently used within some specified distance. The set of cells that interfere with a given cell is called the *interference region* of that cell. In Figure 1, the center cell's interference region consists of all cells within two cell diameters. For certain regular topologies (including those of interest), there exists a smaller set of cells called a *cluster* such that the reuse constraint can be satisfied only if the total number of calls active within each cluster does not exceed the total number of system channels. In Figure 1, one such cluster consists of the center cell and all cells within one cell diameter.

The existing cellular system was proposed by Schulte [1] in 1960. Called *fixed channel allocation* (FCA), it partitions the available spectrum into channel sets (**A**-**G** in Figure 1). The reuse distance constraint is satisfied by assigning these channel sets to the cells in each cluster in a manner determined by a graph coloring problem (c.f. Hale [2]). A base station is allowed to transmit to and from mobiles in its cell only on channels in its assigned channel set.

In Figure 1, a mobile attempting a new call while in the center cell must obtain a channel in segment **G** of the spectrum. This FCA policy clearly is sufficient to insure that no other mobile uses the same channel within the reuse distance, since channels in **G** can not be reused within the interference region of the center cell. However, the policy is not necessary to guarantee that the reuse constraint is satisfied. To demonstrate this, suppose that all channels in **G** are currently used in the center cell, that no other mobiles are placing calls within the interference region in Figure 1, and that one additional mobile in the center cell wishes to initiate a call. Under FCA, this mobile must use a channel in segment **G**, and hence its attempt would be blocked. However, a channel could be found in (for instance) segment **A** that is not used anywhere in the interference region of the center cell. Assigning this channel to the new call would thus not violate any reuse constraint. The realization that FCA is overly restrictive has inspired all other channel allocation policies.

In the late 1960s and early 1970s, a number of alternatives were suggested. Araki [3] introduced the original *dynamic channel allocation* (DCA) policy, which assigns to a new call any channel that is unused in the originating cell's interference region. Cox [4, 5, 6, 7] introduced the concepts of keeping channels in an order (*channel ordering*), assigning channels based on information about channel usage just outside the interference region (*channel assignment*), and reassigning existing calls when a call completes to maintain good channel usage (*channel reassignment*). Engel [8] introduced the concept of initially assigning channels using the FCA policy, but then allowing a base station to borrow a channel from a neighboring cell if it has none available (*channel borrowing*). A plethora of policies have followed from the late 1970s through today. These dynamic channel allocation schemes involve various combinations of permanent channel assignment, channel borrowing, shared pools of channels, channel ordering, channel reassignment, and dynamic adjustment of parameters.

The performance of a given policy has usually been measured by blocking probabilities, under the assumption that call attempts are uniformly distributed among all cells in the system (c.f. [9, 10, 11]). This assumption, however, is increasingly inaccurate in evolving wireless systems with high user mobility. Furthermore, dynamic channel assignment policies reap their greatest gains when the number of calls in each cell are *not* equal. Performance comparisons of channel assignment policies under nonhomogeneous call distributions, therefore, are critical. Such comparisons, however, have been rare. Recently, one study provided capacity bounds for a single channel system under arbitrary arrivals [12]. Previous studies of nonhomogeneous call distributions, however, have generally either assumed a fixed call configuration (c.f. [13, 14]) or a specific set of cell loads (see e.g. [15, 16, 17, 18, 19]). Any comparison with other policies, therefore, is restricted to the particular call configuration or cell loads used. We believe that additional formal metrics might help in the measurement of the performance of cellular channel assignment policies under nonhomogeneous call distributions.

In this paper, we introduce two such metrics to measure the worst-case performance of a channel allocation policy P:

- $CH(S_N, P)$, the worst-case number of channels required to accommodate any configuration of at most N calls in each cell cluster;
- $SS(M, P)$, the set of cell states that can be accommodated using a total of M channels.

The first metric CH provides information about the capability of a channel assignment policy P to accommodate mobility, since the worst-case usually occurs with a skewed call distribution. The second metric SS provides more detailed information as to which call distributions a policy P can accommodate.

The remainder of the paper is organized as follows. In Section 2, we formally define our metrics and briefly review known upper and lower bounds on channel allocation policies under these two metrics. These bounds are generally derived from two extreme policies, fixed channel allocation (FCA) [1] and maximum packing (MP) [9]. In Section 3, we prove a new general lower bound, under the CH metric, on the performance of any channel assignment policy. Section 4 introduces three intermediate channel assignment policies that use well-known ideas such as channel ordering, hybrid assignment, and partitioning. Finally, in Section 5, these three policies are used to demonstrate the variation of performance with the complexity of the channel allocation policy.

2 Known Bounds

2.1 The CH and SS Metrics

First, we define the cell geometry with which we will be working. We assume an infinite regular hexagonal cellular array. Cochannel interference constraints require that a channel not be reused within some minimum distance. This constraint is usually satisfied by specifying an *interference region* around each cell, and requiring that any channel used within a cell not be reused within that cell's interference region. In this paper, we assume that the interference region of cell i is specified as the union of those cells whose centers are a distance of up to and including $r-1$ (a nonnegative integer) from cell i's center, where the distance between centers of neighboring cells is normalized to 1.

Using hexagonal cellular geometry, it can be shown that the minimum distance from the center of cell i to the center of a cell outside cell i's interference region is given by $d = \sqrt{k^2 + kl + l^2}$, where $k = l = r/2$ for r even and $k = (r-1)/2, l = (r+1)/2$ for r odd. Furthermore, a *cluster* can be defined, containing exactly $C = d^2$ cells, as any maximal mutually interfering group of cells. Cell i's interference region can then be represented as the union of all clusters containing cell i. In Figure 1, the interference region is given by $r = 3$ and the cluster size is given by $d^2 = 7$.

Label the cells as $i = 1, 2, \ldots$ and the clusters as $j = 1, 2, \ldots$. Denote the number of calls in cell i as x_i, and the collection of such variables for the entire system as the vector $x = (x_1, x_2, \ldots)$. We denote the number of calls in cluster j by

$$y_j = \sum_{\text{all cells } i \text{ in cluster } j} x_i,$$

and the maximum number of calls in any cluster in the cellular system by

$$\rho(x) = \max\{y_j \mid j \text{ a cluster}\}.$$

Let P be a channel allocation policy. Then we define $CH(x, P)$ to be the number of channels required by policy P to accommodate the call configuration x.

It is well known (c.f. [13]) that $\rho(x)$ is a lower bound on the number of channels required in the cellular system — that is,

$$CH(x, P) \geq \rho(x) \quad \text{for all } P. \tag{1}$$

Furthermore some call configurations have been found for particular cellular topologies and reuse distances which *no* channel assignment policy can accommodate with only $\rho(x)$ channels [13] — that is,

$$\text{For some topologies, there exists an } x \text{ such that } CH(x, P) > \rho(x) \text{ for all } P. \tag{2}$$

Additional bounds have been provided for particular call configurations by representing the channel assignment task as a graph coloring problem [2] or by using additional constraints limiting adjacent channel interference [13].

We are interested, however, not only in how a channel assignment policy accommodates a single call configuration x, but also in how it accommodates a collection of call configurations. In particular, we would like to know how a policy reacts to significant user mobility. We thus define the set of cellular system states in which each cluster carries no more than N calls, in any combination as

$$S_N = \{x \mid y_j \leq N \text{ for all } j\}.$$

Our first metric for measuring the performance of a channel assignment policy P is $CH(S_N, P)$ — the *number of channels* required to accommodate *all cell configurations* in S_N. This metric provides a worst-case measure of the capability of a policy to accommodate mobility. Our second metric is $SS(M, P)$ — the set of cell configurations, or *state space*, that can be accommodated under a particular channel assignment policy P given a total of M channels. $SS(M,P)$ thus displays exactly what types of mobility a channel assignment policy can handle. This second metric is more detailed because given the sets $SS(M,P)$ for all values M and some policy P, we have that

$$CH(S_N, P) = \min\{M \mid S_N \subseteq SS(M,P)\}.$$

We suggest that these two metrics be used in conjunction with more traditional metrics. Probability of blocking provides a measure of the performance of a dynamic channel allocation policy under a particular set of cell loads. The CH and SS metrics provide complementary information about the ability of a policy to accommodate mobility. State space analysis has enjoyed a long tradition in the telecommunications performance community. Calculation of blocking probabilities for dynamic channel allocation policies has usually been accomplished through event-driven simulation, without explicit determination of the policy's achievable state space. Worst-case analysis can not replace such load-specific analysis, but we believe that it can complement it by providing information about bottlenecks, through the CH metric, and information about achievable mobility, through the SS metric.

2.2 Bounds for FCA and MP

The *fixed channel allocation* policy (FCA) for a cellular system with M channels partitions these channels among the C cells in each cluster. It therefore accepts a call to cell i if and only if doing so would result in a state x such that $x_i \leq M/C$. In other words,

$$SS(M, \mathrm{FCA}) = \{x \mid x_i \leq M/C \text{ for all } i\}. \tag{3}$$

It therefore requires as many as CN channels to accommodate any combination of N calls in a cluster (since those calls might all occur in a single cell), namely $CH(S_N, \mathrm{FCA}) = CN$. Since all other proposed channel assignment policies are more complicated than FCA, we think of FCA as providing an upper bound on $CH(S_N, P)$ for reasonable policies P.

On the other hand, the *maximum packing* channel assignment policy (MP) for a cellular system with M channels accepts a call to cell i if and only if doing so would result in a state x such that $y_j \leq M$ for all clusters j containing i. Therefore $SS(M, \mathrm{MP}) = S_M$, and MP only requires N channels to accommodate any combination of N calls in a cluster, namely $CH(S_N, \mathrm{MP}) = N$. (It will be shown in Section 3, however, that maximum packing is unrealizable, and hence only an ideal bound.)

Summarizing, we know that:

$$N \leq CH(S_N, P) \leq CN \tag{4}$$

and

$$SS(M, P) \subseteq S_M \tag{5}$$

for all policies P.[1]

Our goals in the remaining sections of this paper are to provide tighter bounds and to demonstrate the tradeoffs between the performance and the complexity of channel assignment policies under these metrics.

3 A New Lower Bound on the Number of Channels

Individual call configurations x have been presented in the literature for particular cellular topologies and reuse distances that satisfy $x \in S_N$ and yet violate frequency reuse constraints under any channel assignment policy that uses only N channels. In this section, we introduce a family of cell configurations that provides a general lower bound on $CH(S_N, P)$ that is strictly greater than N for all reuse distances.

The configuration includes a simple cycle of adjacent cells, each with N/r calls. Such a configuration is displayed in Figure 2 for the case of $r = 2$. The cycle is constructed so that every connected set of r cells interferes, but no subset of $r+1$ cells interferes. It is therefore clear from expression (1) that the configuration will require a minimum of N channels.

The cycle length l, however, is chosen *not* to be a multiple of the reuse distance r. In Figure 2, the cycle length is chosen to be $l = 9$ and the number of channels $N = 6$. Thus each cell contains $N/r = 3$ calls. Without loss of generality, suppose that channels 1, 2, and 3 are assigned to the cell 0 in the cycle. We must then assign channels 4, 5, 6 to cell 1, since cells 0 and 1 interfere. Similarly, we must assign channels 1, 2, and 3 to cell 2. Proceeding around the cycle clockwise, we will eventually assign channels 4, 5, 6 to cell 7. This leaves cell 8 with no channels that are unused in its interference region. A more detailed analysis of this configuration can furthermore show that no assignment that obeys the reuse constraints is possible with fewer than 8 channels.

The following lemma states that such a cycle can be always be found for any reuse distance $r \geq 2$ and number of channels N.

[1] In the upper bound on CH in expression (4), we ignore policies that perform worse than FCA.

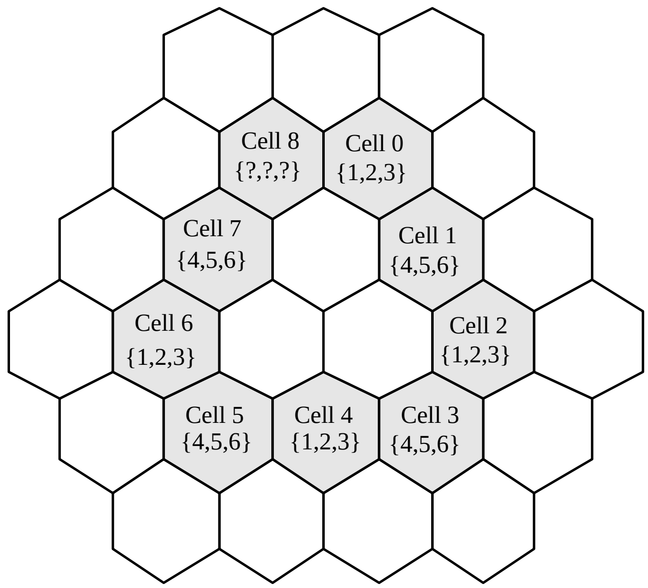

Figure 2: A cycle of cells with 3 calls each.

Lemma 1
For any reuse distance r, there is a cycle of l pairwise adjacent cells, $6r-6 \leq l \leq 6r-3$, such that l is not a multiple of r and no set of $r+1$ cells in the cycle is mutually interfering.

Proof
For all r except $r = 2, 3, 6$, the conditions of the Lemma are satisfied by a regular hexagonal cycle of $l = 6(r-1) = 6r-6$ cells. It is easy to verify that such a cycle satisfies the interference conditions for all r. However, for $r = 2, 3, 6$, it has length l that is a multiple of r. We address these three cases separately:

- $r = 6$: We can elongate two opposite sides of the regular hexagonal cycle by a single cell each, so that two opposite sides of the cycle contain seven cells each, and the other four side contain six cells each. This yields a cycle of length $l = 6r-4 = 32$ that satisfies the interference conditions and whose length is not a multiple of r.

- $r = 3$: Again, we elongate two opposite sides of the regular hexagonal cycle by a single cell each, so that two opposite sides of the cycle contain four cells each, and the other four side contain three cells each. This yields a cycle of length $l = 6r-4 = 14$ that satisfies the interference conditions and whose length is not a multiple of r.

- $r = 2$: A hexagonal cycle of nine cells whose sides alternate between having two and three cells satisfies the interference conditions and has length $l = 6r - 3 = 9$ that is not a multiple of r. (This is the cycle pictured in Figure 2.)

This yields for every r a cycle of length between $6r - 6$ and $6r - 3$ that satisfies the conditions of the Lemma. □

The following theorem uses this cycle to prove a formal lower bound on the worst-case number of channels required to satisfy a configuration of calls in cells with at most N calls per cluster.

Theorem 2
Suppose we have a cycle of pairwise adjacent cells of length l, where l is not a multiple of r and no set of $r + 1$ cells in the cycle is mutually interfering. Then we can assign calls to each cell in the cycle in such a way that there are at most N calls per cluster but at least $N + \left\lceil \frac{\lfloor N/r \rfloor}{\lfloor l/r \rfloor} \right\rceil$ channels are needed to satisfy the calls within the reuse constraints.

Proof
We will prove the theorem for the case where N is a multiple of r; the proof of the general case is nearly identical.

Label the cells of the cycle $i_0, i_1, \ldots, i_{l-1}$, and assign exactly N/r calls to each cell in the cycle. Note that this assigns at most N calls to each cluster. Assume that we can satisfy this configuration of calls with $N + k$ channels. Since cells i_0 through i_{r-1} of the cycle are mutually interfering, they together must use N channels. Therefore cells i_0 and i_r must have at least $N/r - k$ of their assigned channels in common. Similarly, cells i_r and i_{2r} must have least $N/r - k$ of their assigned channels in common — it follows that i_0 and i_{2r} must have at least $N/r - 2k$ of their assigned channels in common. It is straightforward to verify that for any $q \leq \lfloor l/r \rfloor$, cells i_0 and i_{qr} must have at least $N/r - qk$ of their assigned channels in common. Consider cell $i_{\lfloor l/r \rfloor r}$. This cell must have at least $N/r - \lfloor l/r \rfloor k$ of its assigned channels in common with i_0. However, these two cells are within the reuse distance r of each other, so they cannot have any channels in common. Therefore it must be the case that $N/r - \lfloor l/r \rfloor k$ is at most zero. It follows that $k \geq \frac{N/r}{\lfloor l/r \rfloor}$, and therefore that the number of channels required to satisfy this configuration is at least $N + \left\lceil \frac{N/r}{\lfloor l/r \rfloor} \right\rceil = N + \left\lceil \frac{\lfloor N/r \rfloor}{\lfloor l/r \rfloor} \right\rceil$.

When N is not a multiple of r, assigning either $\lfloor N/r \rfloor$ or $\lceil N/r \rceil$ calls to each cell in the cycle in an appropriate fashion will yield the general form of the theorem by an identical argument. □

Applying Theorem 2 to the cycles whose existence was proved in the previous lemma yields the following corollary:

Corollary 3
For any reuse distance r, there is a configuration of calls in cells with at most N calls per cluster for which the number of channels required to satisfy the calls is $N + \lceil \lfloor N/r \rfloor /4 \rceil$ if $r \leq 5$, and $N + \lceil \lfloor N/r \rfloor /5 \rceil$ if $r \geq 6$. □

Note that this lower bound is strictly greater than the previous lower bound, $\rho(x)$ in expression (2), with the trivial requirement that $N \geq r$, since here $\rho(x) = N$. Furthermore, this bound applies to all channel assignment policies P and all reuse distances r, using only cochannel constraints. Maximum packing, therefore, is an unrealizable channel assignment policy, and represents only an unachievable performance bound. It follows that the lower bound of N in expression (4) is a strict inequality. Furthermore it also follows that no policy can accommodate all combinations of M calls in a cluster with only M channels, namely $SS(M, P) \subset S_M$ for all policies P.

4 Three Intermediate Channel Allocation Policies

In this section, we introduce three intermediate channel assignment policies, based on commonly used ideas of channel ordering, hybrid assignment, and partitioning. These policies were chosen due to the prevalent usage of their underlying concepts in the dynamic channel assignment literature and to our ability to construct analyzable variants of them. Furthermore, little is known about the performance benefits of these concepts under nonhomogeneous loads [5, 16]. These policies will be used to demonstrate the variation of performance with the complexity of the channel allocation policy, and are not meant to outperform other policies proposed in the literature.

4.1 Circular Ordering

Our first scheme, Circular Ordering (CO), is relatively simple and based on channel ordering, which has been widely used in dynamic channel allocation policies since it was first suggested by Cox [5].
Circular Ordering: Each cell is marked as in FCA with one of C letters (**A**, **B**, **C**, ...), such that cells marked with the same letter are noninterfering. This marking is shown in Figure 3(a) for a system with $r = 2$. The M channels are conceptually arranged on a circle and each of the C segments is assigned a *center*;

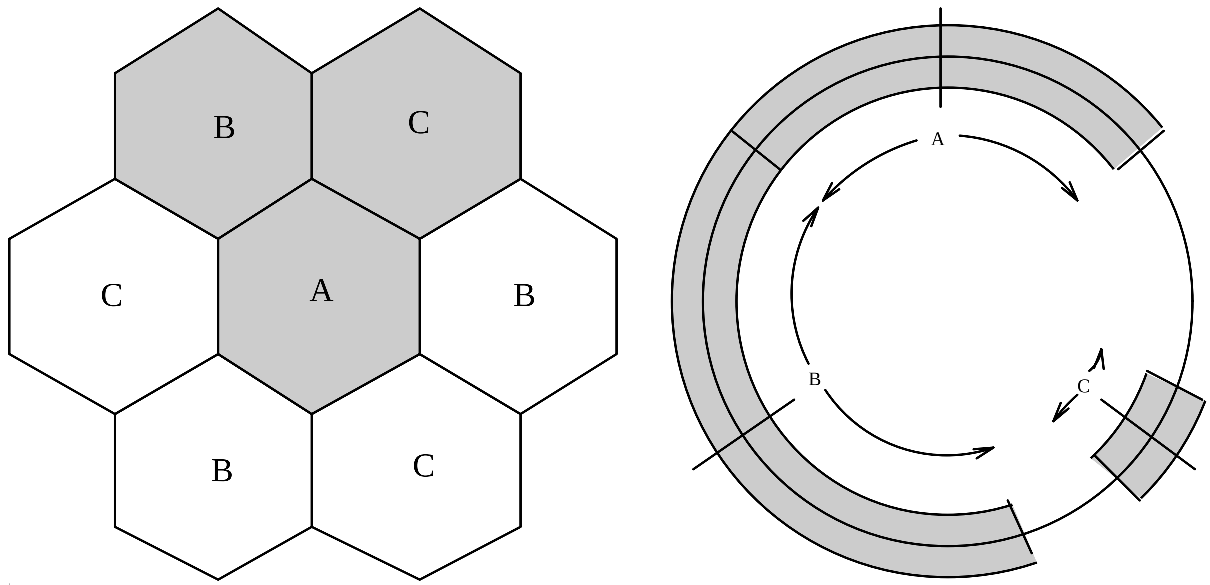

Figure 3: (a) A cell's clusters. (b) A channel assignment circle for one cluster.

these centers are equally spaced around the circle. For each cluster, this circle is used to track which channels are currently in use in each cell. In Figure 3(a), there are six clusters containing the center cell, one of which is shaded. A representative circle is shown in Figure 3(b).

When a mobile requests a channel from cell i, the cell attempts to assign it the channel, among those unused by cell i and between adjacent centers on the circle, that is closest to cell i's center. If no such channel exists or if this channel is occupied by another cell in any cluster containing cell i, then the call request is blocked. When any channel is released, the cell reassigns that channel to the call currently occupying the channel furthest from the cell's center. Thus an assignment is maintained in which the calls in each cell i occupy a contiguous range of x_i channels,

$$\left(k\frac{N}{2} - \left\lceil \frac{x_i - 1}{2} \right\rceil\right) \bmod M, \quad \ldots, \quad \left(k\frac{N}{2} + \left\lfloor \frac{x_i - 1}{2} \right\rfloor\right) \bmod M,$$

symmetric about the cell's center, where $k = 1$ for type **A** cells, $k = 2$ for type **B** cells, etc. □

In Figure 3(b), the shaded portions represent channels occupied by cells **A**, **B**, and **C** in the shaded cluster in Figure 3(a). For this example, a new call in the shaded cell **A** or **B** would be blocked. A new call in the shaded cell **C** would be assigned to the unshaded channel closest to its center, if this channel were also free in the circles corresponding to the other clusters containing that cell.

The states that are achievable under this policy are described in the following theorem.

Theorem 4
Circular Ordering, with M channels (assumed to be a multiple of the cluster size C), can satisfy any state in S_M such that:

$$x_{i_1}+x_{i_2} \leq \frac{2M}{C} \quad \text{for all cells } i_1,\ i_2 \text{ in a common cluster and adjacent on the circle.} \tag{6}$$

Thus $SS(M, \mathrm{CO})$ = those configurations satisfying expression (6) $\subset S_M$.

Furthermore, Circular Ordering can satisfy any configuration of calls in cells with at most N calls per cluster using $\frac{CN}{2}$ channels when CN is even, $\lceil \frac{CN}{2} \rceil = \frac{CN+1}{2}$ channels otherwise. Thus $CC(S_N, \mathrm{CO}) = \lceil \frac{CN}{2} \rceil$.

Proof
(We will give the proof for the case where N is even. The proofs for when N is odd but C is even, and for when CN is odd are nearly identical.)

Let $M = \frac{CN}{2}$. We show that if each cluster contains at most N calls, then no two of the ranges of channels assigned to two cells in the same cluster overlap. The Theorem follows.

Suppose that each cluster contains at most N calls. Let i_1 and i_2 be two cells in the same cluster, of types k_{i_1} and k_{i_2} and with x_{i_1} and x_{i_2} calls respectively; then $x_{i_1} + x_{i_2} \leq N$. Without loss of generality, assume that $k_{i_1} < k_{i_2}$. The calls in each cell are assigned channels

$$\left(k_{i_1}\frac{N}{2} - \left\lceil \frac{x_{i_1}-1}{2} \right\rceil\right) \bmod \frac{CN}{2},\ \ \ldots,\ \ \left(k_{i_1}\frac{N}{2} + \left\lfloor \frac{x_{i_1}-1}{2} \right\rfloor\right) \bmod \frac{CN}{2}$$

and

$$\left(k_{i_2}\frac{N}{2} - \left\lceil \frac{x_{i_2}-1}{2} \right\rceil\right) \bmod \frac{CN}{2},\ \ \ldots,\ \ \left(k_{i_2}\frac{N}{2} + \left\lfloor \frac{x_{i_2}-1}{2} \right\rfloor\right) \bmod \frac{CN}{2}.$$

The high end of i_1's range overlaps the low end of i_2's range only if

$$k_{i_1}\frac{N}{2} + \left\lfloor \frac{x_{i_1}-1}{2} \right\rfloor \geq k_{i_2}\frac{N}{2} - \left\lceil \frac{x_{i_2}-1}{2} \right\rceil,$$

or equivalently,

$$(k_{i_2} - k_{i_1})\frac{N}{2} \leq \left\lfloor \frac{x_{i_1}-1}{2} \right\rfloor + \left\lceil \frac{x_{i_2}-1}{2} \right\rceil.$$

However, since $x_{i_1} + x_{i_2} \leq N$, we have

$$\left\lfloor \frac{x_{i_1} - 1}{2} \right\rfloor + \left\lceil \frac{x_{i_2} - 1}{2} \right\rceil \leq \frac{x_{i_1} - 1}{2} + \frac{x_{i_2} - 1}{2} + \frac{1}{2} \leq \frac{N}{2} - \frac{1}{2},$$

but from the fact that $k_{i_2} - k_{i_1} \geq 1$,

$$(k_{i_2} - k_{i_1})\frac{N}{2} \geq \frac{N}{2},$$

so it can never be the case that the high end of i_1's range of channels and the low end of i_2's range of channels overlap. A similar argument shows that the low end of i_1's range and the high end of i_2's range of channels cannot overlap either. We conclude that the ranges of channels assigned to the calls in cells i_1 and i_2 are disjoint. Since i_1 and i_2 were arbitrary interfering cells, and the assignment of calls to cells was arbitrary, the Theorem follows. □

The modifications to this proof required to establish the more general result are minimal. When N is odd, rather than using the C centers to divide each circle of $\frac{CN}{2}$ channels into C equal arcs of $\frac{N}{2}$ channels, we instead divide it into C alternating arcs of $\lceil \frac{N}{2} \rceil$ and $\lfloor \frac{N}{2} \rfloor$ channels. When CN is odd the fact that there is one more arc of size $\lceil \frac{N}{2} \rceil$ than of size $\lfloor \frac{N}{2} \rfloor$ leads to the increase in the bound.

Note that this scheme allows more sharing than FCA since FCA's state space, as given by expression (3), is a strict subset of Circular Ordering's state space, as given by expression (6). Those states in $SS(M, \text{CO})$ but not in $SS(M, \text{FCA})$ represent call configurations CO can handle but FCA can not. This sharing results in a reduction by one half of the number of required channels to accommodate any configuration of N calls in a cluster. Circular Ordering thus achieves a performance, with respect to our two metrics, that is near the midpoint of the ranges defined by FCA and MP in expressions (4) and (5). This increase in performance above that achieved by FCA directly measures Circular Ordering's superior ability to react to mobility, and helps explain the popularity of channel ordering in the literature.

4.2 Hybrid Graph Coloring

Better performance in our metrics can be achieved using more sharing. Hybrid Graph Coloring (HGC) is a bit more complex and is based on hybrid allocation [16]. We define $M = \frac{C+\frac{7}{2}}{3}N$, and we assume for convenience that N is an integer multiple of 6, so that M is an integer multiple of $2C + 7$. The general version is considered in Theorem 5.

Hybrid Graph Coloring: Each cell is marked as in FCA with one of C letters $(\mathbf{A}, \mathbf{B}, \mathbf{C}, \ldots)$, such that cells marked with the same letter are noninterfering. This marking is shown in Figure 4 for a system with $r = 2$. A total of $CN/3$ channels

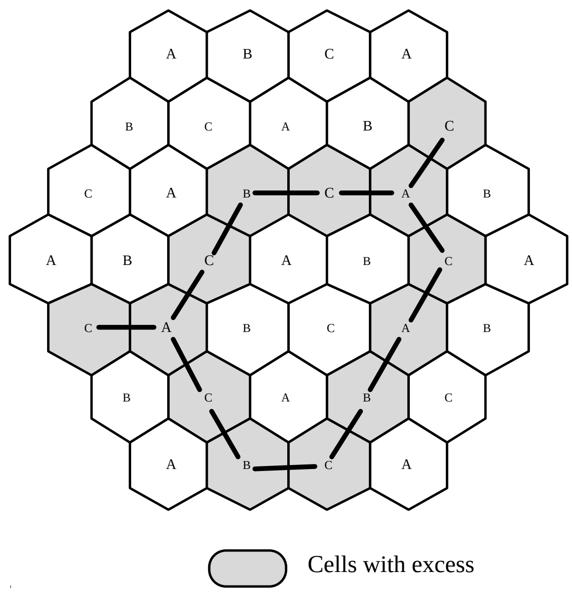

Figure 4: Cell marking and construction of the excess graph.

are partitioned equally among the C markings. This leaves $M - CN/3 = 7N/6$ channels to be shared among all cells.

When a mobile requests a channel from cell i, the cell first attempts to assign it a channel in its segment of the partition. If no such channel is available, a graph is constructed with a node for each cell carrying greater than $N/3$ calls (including the new call) and with an edge between any two nodes that interfere. Denote all calls not assigned channels in the partition as *excess* calls, and denote the corresponding graph as the *excess graph* for the configuration. Such a graph is shown in Figure 4. Let h be the number of colors required to color the excess graph's nodes so that no adjacent nodes receive the same color. Use Circular Ordering, with the remaining $M - CN/3 = 7N/6$ channels and h centers, to satisfy the excess calls, if possible. If Circular Ordering can not satisfy all excess calls, block the new call.

When a mobile terminates a call, if that call was assigned a channel by Circular Ordering in the excess graph, proceed as in Circular Ordering. If the call was assigned one of the fixed channels for its cell, terminate one of calls in the excess graph node for the cell and assign to it the fixed channel that was just released. □

The states that are achievable under this policy are described in the following theorem.

Theorem 5
Hybrid Graph Coloring, with M channels (assumed to be a multiple of $2C+7$), can accommodate a state space bounded by:

$$SS(M, \mathrm{CO}) \subset SS(M, \mathrm{HGC}) \subseteq S_M \cap \{x_i \leq \frac{9M}{2C+7} \quad \forall i\}. \tag{7}$$

Thus $SS(M, \mathrm{CO}) \subset SS(M, \mathrm{HGC}) \subset S_M$.

Furthermore, Hybrid Graph Coloring can satisfy any configuration of calls in cells with at most N calls per cluster using $\frac{C+7/2}{3}N$ channels whenever N is a multiple of six, and $(C+7/2)\lfloor\frac{N}{3}\rfloor+14$ channels otherwise. Thus $CH(S_N, \mathrm{HGC}) \leq \frac{C+7/2}{3}N+14$.

Proof
Due to the variability of h, the number of colors required, we do not have a exact characterization of Hybrid Graph Coloring's state space. The number of shared channels is $M - \frac{CN}{3} = 7N/6$. The maximum number of channels any one cell can occupy is therefore $N/3 + 7N/6 = 3N/2 = \frac{9M}{2C+7}$. The upper bound follows.

For the lower bound, compare $SS(M, \mathrm{HGC})$ to $SS(M, \mathrm{CO})$. Assume a configuration $x \in SS(M, \mathrm{CO})$. Use the coloring corresponding the Circular Ordering's markings to label the excess graph. Then any pair of cells x_{i1} and x_{i2} in the excess graph that are also adjacent on the ordering circle have an excess of at most $(x_{i1} - \frac{N}{3}) + (x_{i2} - \frac{N}{3})$. Since $x_{i_1} + x_{i_2} \leq \frac{2M}{C}$, from expression (6), their joint excess is at most $\frac{2M}{C} - \frac{2N}{3}$. On the other hand, Hybrid Graph Coloring can accommodate these cells if their joint excess is at most $\frac{2}{C}$ times the number of shared channels, $7N/6$. The two quantities are equal; it follows that $x \in SS(M, \mathrm{HGC})$ and hence $SS(M, \mathrm{CO}) \subseteq SS(M, \mathrm{HGC})$. In addition, the state $(\frac{9M}{2C+7}, 0, 0, \ldots)$ can be accommodated under HGC but not under CO, since $\frac{9M}{2C+7} > \frac{2M}{C}$ for $C > 2$. Hence $SS(M, \mathrm{CO}) \subset SS(M, \mathrm{HGC})$, and the lower bound is established.

It is clear that Hybrid Graph Coloring never assigns the same channel to two calls in the same cluster; all that remains to be shown is that the claimed bound on the number of channels used is correct. For clarity of proof, we will first assume that N is a multiple of six.

If a pair of cells in the same cluster have excess, and the cluster carries at most N calls, then the total excess of that pair can be at most $N/3$. Therefore Circular Ordering can satisfy the excess calls with $(h/2)(N/3)$ additional channels. Let g be the maximum degree of the excess graph. Then the graph can be colored with $g+1$ colors, so $h \leq g+1$. Circular Ordering therefore will require at most $\frac{(g+1)N}{6}$ additional channels. Since satisfying up to the first $N/3$ calls in each cell is accomplished using $CN/3$ channels, and the number of shared channels available is $M - CN/3 = 7N/6$, it will suffice to show that the maximum degree of the excess graph is six.

Consider any node in the excess graph; this node corresponds to a cell i that initially had excess after $N/3$ of its calls were assigned channels. Consider the region consisting of all cells within the reuse distance of cell i. This region can be covered by six clusters, all containing i, in such a way that every cell in the region is in at least one cluster. Since at most two cells in each cluster can have excess calls, and i is already in each cluster, each cluster can contain at most one cell other than i with excess. It follows that no node in the excess graph has more than six neighbors, so that the maximum degree of the excess graph is six. Therefore Hybrid Graph Coloring requires a total of $CN/3 + 7N/6 = \frac{C+7/2}{3}N$ channels.

For the case in which N is not a multiple of six, we modify the scheme to initially satisfy up to $\lfloor N/3 \rfloor$ calls in each cell, so that any two cells with excess in the same cluster must have joint excess at most $\lfloor N/3 \rfloor + 2$. An identical analysis to the one above yields the general form of the theorem. □

We note that for the case $C = 3$, the lower degree of the resulting excess graph together with its planarity can be used to improve the number of channels to $\frac{C+1}{3}N = \frac{4}{3}N$ when N is a multiple of three.

Hybrid Graph Coloring allows more sharing than Circular Ordering, since $SS(M, \mathrm{CO}) \subset SS(M, \mathrm{HGC})$. This additional sharing results in a reduction in the number of required channels to accommodate any N calls in a cluster, from $\frac{CN}{2}(1 + o(1))$ for Circular Ordering to $\frac{CN}{3}(1 + o(1))$ for Hybrid Graph Coloring. The increase in performance under both metrics directly demonstrates the advantage in using hybrid policies to react to mobility and helps explain their use in many channel assignment policies. This performance gain, however, is achieved at the cost of additional complexity. The allocation of a channel in Hybrid Graph Coloring may rely on the construction of a graph of arbitrary size, and hence this policy is likely to be more complicated than Circular Ordering.

The performance can be further improved by generalizing Hybrid Graph Coloring. If the number of channels assigned to each cell in the first stage of Hybrid Graph Coloring is reduced, then the worst-case number of channels required can also be reduced. In particular, for any integer k, we can choose to satisfy only up to N/k calls initially in each cell. The resulting theorem is as follows (stated without proof):

Theorem 6

For any integer $k \geq 2$, a modified version of Hybrid Graph Coloring (denoted by HGC'*) can satisfy any configuration of calls in cells where there are at most N calls per cluster with $(\frac{C}{k} + 3k - 6)N$ channels when N is a multiple of $2k$, $(C - 6k)\lfloor \frac{N}{k} \rfloor + 3kN + 1$ channels otherwise. Thus $CH(S_N, \mathrm{HGC}') \leq (\frac{C}{k} + 3k - 6)N + 1$. (When $k = 1$,* HGC' *is equivalent to* FCA.*)* □

This increased performance is achieved at the cost of increasing the size of the graph of cells with excess demand. As seen in previous studies of hybrid policies, the number of fixed channels can be chosen according to the expected amount of mobility. Here, the worst-case performance is explicitly represented. For a fixed cluster size C, the optimal choice of k is $\Theta(\sqrt{C})$, and the resulting upper bound on the number of channels required is $\Theta(\sqrt{C}N)$. Recall that $N < CH(S_N, P) \leq CN$ for all reasonable P. Circular Ordering and Hybrid Graph Coloring require $\Omega(CN)$ and $\Omega(\sqrt{C}N)$ channels, respectively.[2] The lower bound, however, still leaves the possibility of policies that require only $O(N)$ channels.

4.3 Cluster Partitioning

One policy P that achieves $CH(S_N, P) = O(N)$, which we call Cluster Partitioning (CLP), can be constructed by using the concept of partitioning, but among clusters, not cells.

Cluster Partitioning: The cellular array is partitioned into clusters of size C, as in FCA. Each *cluster*, not cell, is assigned a group of channels marked (**A**, **B**, **C**, ...) in such a way that any pair of clusters containing cells that are within distance $r-1$ are assigned different markings. A channel in any cluster's channel group can be assigned to any one call in that cluster, arbitrarily. □

Such a marking is shown in Figure 5 for a system with $r = 2$. It can be shown that for any cluster size C, only four channel groups are required to insure that any two cells using the same channel (in two different clusters) are noninterfering. The states that are achievable under this policy are described in the following theorem.

Theorem 7

Cluster Partitioning, with M channels (assumed to be a multiple of 4), can satisfy any state such that

$$y_j \leq M/4 \text{ for all nonoverlapping clusters } j. \tag{8}$$

Thus $SS(M, \mathrm{CLP})$ = the set of all configurations that satisfy expression (8).

Furthermore, Cluster Partitioning can satisfy any configuration of calls in cells with at most N calls per cluster using $4N$ channels. Thus $CH(S_N, \mathrm{CLP}) = 4N$.

Proof

For a cluster i, we define the *snowflake centered at i* to be the union of all clusters j in the partition such that there is a cell in cluster i and a cell in cluster j that

[2] Roughly speaking, just as a quantity is said to be $O(f(N))$ if its asymptotic rate of growth is *at most* $c \cdot f(N)$ for some constant c, a quantity is said to be $\Omega(f(N))$ if its asymptotic rate of growth is *at least* $c \cdot f(N)$ for some constant c. A quantity that is both $O(f(N))$ and $\Omega(f(N))$ is said to be $\Theta(f(N))$. For a more detailed discussion, see [20].

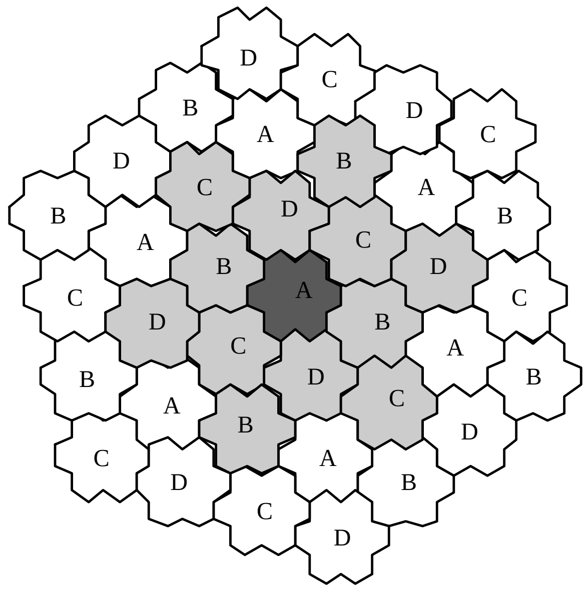

Figure 5: Partitioning channel groups among clusters.

are within distance $r-1$ of each other (and therefore are mutually interfering).

For each cluster i, the snowflake centered at i consists of thirteen clusters:

1. The cluster i itself;
2. The six clusters that are adjacent to i in the partition;
3. The six additional clusters that are adjacent to two of the neighboring clusters of i.

It is straightforward to verify that any cell that is within distance $r-1$ of a cell in cluster i is contained in one of the clusters in the snowflake centered at i.

Furthermore, it is always possible to four-color the clusters in the partition such that for every cluster i, all the clusters in the snowflake centered around i are assigned different colors than i is. Such a four-coloring is illustrated in Figure 5 for $r = 3$. Therefore any two cells within distance $r-1$ of each other are in clusters that are assigned different colors.

Any two calls in cells within distance $r-1$ of each other will be assigned different channels, as follows: In they are in the same cluster, they will be assigned distinct channels from the same set of N channels. If they are in different clusters, then they will be assigned channels from two disjoint sets of N channels, since the clusters containing the two cells were assigned different colors.

It follows that no two calls in cells within distance $r-1$ of each other are assigned the same channel, and therefore that any configuration of calls in cells can be satisfied with $4N$ channels. □

Cluster Partitioning thus achieves a worst-case number of required channels of $4N$ to accommodate any configuration of at most N calls in a cluster, compared to $\Omega(\sqrt{C}N)$ for previous schemes. It does so, however, by severely restricting sharing of channels. Its state space does *not* contain $SS(M, \mathrm{FCA})$, since a configuration consisting of M/C calls in each cell could be accommodated under FCA but not under CLP. CLP is thus tailored to accommodate skewed call configurations using large cluster sizes.

5 Tradeoffs Between Performance and Complexity

It is often helpful to measure a channel assignment policy by the amount of information required to choose a channel, and by the number of reassignments of channels to existing calls required [21]. One or more reassignments can be required at call termination to maintain channel ordering and/or at call setup to accommodate a new call. Increased information or increased reconfigurations represent additional complexity which should be justified by increased performance as demonstrated by such metrics as CH or SS.

FCA represents the simplest policy on both accounts. It requires only knowledge of the channels assigned within the cell in which a call originates. Furthermore, it never reassigns an existing call to another channel, except during handoff. Maximum packing, on the other hand, represents the asymptotic upper limit. It potentially requires knowledge of channel assignments in the entire (here assumed infinite) cellular array to decide whether a new call can be accommodated. MP also potentially requires an infinite number of channel reassignments at call setup to free the required channel.

Circular Ordering attempts to assign a new call a channel that is adjacent to those already used in the cell. The originating cell must check with any other cell in a common cluster that is a neighbor on its channel assignment circle. This requires collection of information within the originating cell's interference region. Furthermore, since Circular Ordering requires that channel assignments be contiguous within each cell, it may require one channel reassignment at call termination. The additional complexity above that required for FCA, in both knowledge and reconfigurations, pays off with higher performance, as measured by the two metrics discussed above. These results provide evidence for the ability of channel ordering to accommodate mobility.

Hybrid Graph Coloring first attempts to assign a new call a channel from its permanent channel group. This requires no information outside the originating cell. If all such channels are occupied, however, then that cell must enter the excess graph. This graph must be colored to assign markings to each cell in the graph, and then Circular Ordering must be applied to assign channels to each cell. The graph coloring requires knowledge from channel assignments in each cell in the excess graph, and this graph can contain an unlimited number of cells. The information required, therefore, is unbounded. Similarly, the number of channel reassignments required at call setup is unbounded. The additional complexity above that required for Circular Ordering, in both knowledge and reconfigurations, pays off with higher performance over Circular Ordering as measured by both metrics discussed above. These results provide evidence for the ability of more complex policies to further accommodate mobility.

Cluster Partitioning assigns a new call a channel from the channel group belonging to that cell's cluster. This requires only information from that cluster. Furthermore, no channel reassignments are required, except for handoffs. Under the metric of accommodating any combination of N calls within a cluster, this method can outperform Hybrid Graph Coloring for large cluster sizes with less complexity. It also outperforms FCA, for cluster sizes exceeding 7, with little additional complexity, under this CH metric. Cluster Partitioning, however, does not necessarily outperform even FCA under the SS metric. Although CLP approaches the lower bound on the worst-case number of channels required for any channel allocation policy, given in expression (4), it therefore does *not* approach the upper bound on the set of achievable states, given in expression (5). These results leave open the question of whether more complicated policies could approach both these bounds.

It should also be noted that the above pairwise comparisons of policies *do not* necessarily hold under other metrics. If the desired loading is known, one would generally prefer schemes that carry a higher average number of calls. In particular, Cluster Partitioning allows little sharing among cells, and would perform badly under even loading. Using average throughput as a performance measurement, however, requires a statistical traffic model and a specific load distribution. In contrast, this study has concentrated on provable bounds for the performance of general channel assignment policies under a range of load configurations corresponding to significant user mobility.

6 Conclusion

We have introduced two performance metrics for channel allocation policies in cellular systems. The first metric, $CH(S_N, P)$, measures the worst-case number of channels required to accommodate N calls in a cell cluster, in any combination, using the policy P. The second more detailed metric, $SS(M, P)$, measures the set of cell states that can be achieved with M channels using the policy P.

We have proven a new general lower bound, under the CH metric, on the performance of any channel assignment policy. This bound demonstrates the degree to which the maximum packing policy is unachievable in cellular systems with different cluster sizes. Three intermediate policies — Circular Ordering, Hybrid Graph Coloring, and Cluster Partitioning — were also introduced that demonstrate how commonly used channel allocation mechanisms achieve a wide range of performances with respect to these metrics. Finally, we recognized that increased performance is usually achieved at the cost of increased complexity, in terms of required information and number of forced reconfigurations.

In the literature, many cellular channel assignment policies have been proposed to improve efficiency above that resulting from fixed channel allocation. The performance of these policies, however, has rarely been compared due to a lack of formal metrics, particularly under nonhomogeneous call distributions. It is our hope these two metrics may be used to gain some insight into the relative performance of various channel assignment policies in systems with significant user mobility.

References

[1] H. Schulte and W. Cornell, Multi-area mobile telephone system, IRE Transactions on Vehicular Communications 9, (1960) 49–53.

[2] W. K. Hale, Frequency assignment: theory and applications, Proceedings of the IEEE 68, (1980) 1497–1514.

[3] K. Araki, Fundamental problems of nation-wide mobile radio telephone system, Review of the Electrical Communication Laboratory (Japan) 16, (1968) 357–373.

[4] D. C. Cox and D. O. Reudink, Dynamic channel assignment in high-capacity mobile communication systems, Bell System Technical Journal 50, (1971) 1833–1857.

[5] D. C. Cox and D. O. Reudink, A comparison of some channel assignment strategies in large-scale mobile communication systems, IEEE Transactions on Communications 20, (1972) 190–195.

[6] D. C. Cox and D. O. Reudink, Dynamic channel assignment in two-dimensional large-scale mobile radio systems, Bell System Technical Journal 51, (1927) 1611–1629.

[7] D. C. Cox and D. O. Reudink, Increasing channel occupancy in large-scale mobile radio systems: dynamic channel reassignment, IEEE Transactions on Vehicular Technology 22, (1973) 218–222.

[8] J. S. Engel and M. M. Peritsky, Statistically-optimum dynamic server assignment in systems with interfering servers, IEEE Transactions on Vehicular Technology 22, (1973) 203–209.

[9] D. Everitt and D. Manfield, Performance analysis of cellular mobile communication systems with dynamic channel assignment, IEEE Journal on Selected Areas in Communications 7, (1989) 1172–1179.

[10] P.-A. Raymond, Performance analysis of cellular networks, IEEE Transactions on Communications 39, (1991) 1787–1793.

[11] S. Jordan and A. Khan, A performance bound on dynamic channel allocation in cellular

systems: equal load, IEEE Transactions on Vehicular Technology 43, (1994) 333–344.

[12] L. J. Cimini, G. J. Foschini, and L. A. Shepp, Single-channel user-capacity calculations for self-organizing cellular systems, IEEE Transactions on Communications 42, (1994) 3137–3143.

[13] A. Gamst, Some lower bounds for a class of frequency assignment problems, IEEE Transactions on Vehicular Technology 35, (1986) 8–14.

[14] L. Cimini, G. Foschini, and C.-L. I, Call blocking performance of distributed algorithms for dynamic channel allocation in microcells, in *Conference Record of the International Conference on Communications (ICC)*, IEEE, 1992, pp.345A.4.1–345A.4.6.

[15] D. C. Cox and D. O. Reudink, Effects of some nonuniform spatial demand profiles on mobile radio system performance, IEEE Transactions on Vehicular Technology 21, (1972) 62–67.

[16] T. J. Kahwa and N. D. Georganas, A hybrid channel assignment scheme in large-scale, cellular-structured mobile communication systems, IEEE Transactions on Communications 26, (1978) 432–438.

[17] M. Zhang and T.-S. Yum, Comparisons of channel-assignment strategies in cellular mobile telephone systems, IEEE Transactions on Vehicular Technology 38, (1989) 211–215.

[18] A. Khan and S. Jordan, A performance bound on dynamic channel allocation in cellular networks: unequal load. Preprint, Feb. 1994.

[19] R. J. McEliece and K. N. Sivarajan, Performance limits for channelized cellular telephone systems, IEEE Transactions on Information Theory 40, (1994) 21–34.

[20] T. Cormen, C. Leiserson, and R. Rivest, *Introduction to Algorithms*, MIT Press/McGraw-Hill, Cambridge, 1990.

[21] S. Jordan, Resource allocation in wireless networks, Journal of High Speed Networks, to appear (1995).