38

Online Call Admission Control in Wireless Cellular Networks

Ioannis Caragiannis
Christos Kaklamanis
Evi Papaioannou
University of Patras

38.1 Introduction

We study frequency spectrum management issues in wireless cellular networks where base stations connected through a high-speed network are used to build the required communication infrastructure. A geographical area in which communication takes place is divided into regions. Each region is the calling area of a base station. In such systems, communication is established in the following way. When a user A wishes to communicate with some other user B, a path must be established between the base stations of the regions where users A and B are located. Then communication is performed in three steps: (i) wireless communication between A and its base station, (ii) communication between the base stations, and (iii) wireless communication between B and its base station. At least one base station is involved in the communication even if both users are located in the same region or only one of the two users is part of the cellular network (and the other uses for example the PSTN). Improving the access of users to base stations is the aim of this work.

The network topology usually adopted [9–11, 21] is the one shown in the left part of Figure 38.1. All regions are regular hexagons (cells) of the same size. This shape results from the assumption of uniform distribution of identical base stations within the network, as well as from the fact that the calling area of a base station is a circle which, for simplicity reasons, is idealized as a regular hexagon. Owing to the shape of the cells, we call these networks wireless cellular networks.

Many users of the same region can communicate simultaneously with their base station through Frequency Division Multiplexing (FDM). The base station is responsible for allocating distinct frequencies from the available spectrum to users so that signal interference is avoided. Signal interference manifests

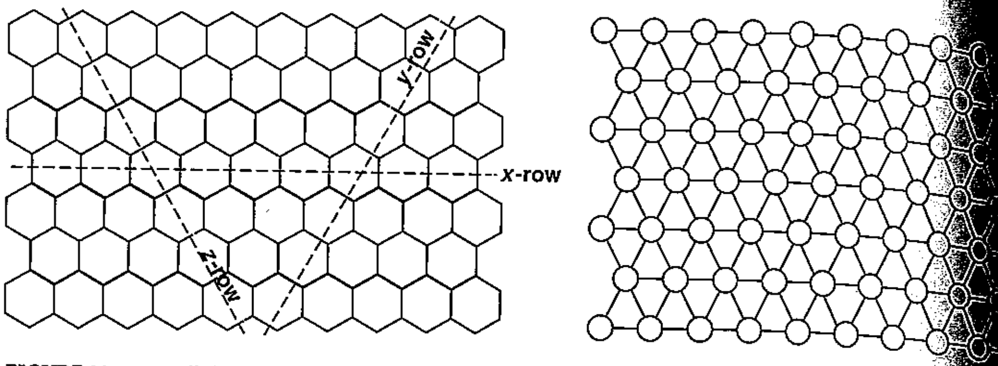

FIGURE 38.1 A cellular network and its interference graph.

itself when the same frequency is assigned to users located in the same or adjacent cells. It c... represented by an *interference graph G* whose vertices correspond to cells, and an edge (u, v) ind... that the assignment of the same frequency to two users lying at the cells corresponding to nodes... will cause signal interference. The interference graph of a cellular network is depicted in the right... Figure 38.1. If the assumption of uniform distribution of identical base stations does not hold, arb... interference graphs can be used to model the underlying network.

Since the spectrum of available frequencies is limited, important engineering problems related... efficient reuse of frequencies arise [12, 13, 17, 18]. We study the *call admission control* (or, simpl... control) problem that is defined as follows: Given users that wish to communicate with their base... the *call control* problem in a network that supports a spectrum of w available frequencies is to... frequencies to users so that at most w frequencies are used in total, signal interference is avoided, a... number of users served is maximized.

We assume that calls corresponding to users that wish to communicate with their base station a... in the cells of the network in an online manner. When a call arrives, a call control algorithm decides... to accept the call (assigning a frequency to it) or to reject it. Once a call is accepted, it cannot be... (preempted). Furthermore, the frequency assigned to the call cannot be changed in the future. We... that all calls have infinite duration; this assumption is equivalent to considering calls of similar du...

We use competitive analysis [3] to evaluate the performance of online call control algorithms wi... competitive ratio being the performance measure. In this setting, given a sequence of calls, the perform... of an online call control algorithm A is compared to the performance of the optimal offline algo... *OPT*. Let $B_A(\sigma)$ be the benefit of the online algorithm A on the sequence of calls σ, that is, the nu... of calls of σ accepted by A, and $B_{\mathrm{OPT}}(\sigma)$ the benefit of the optimal algorithm *OPT*. If A is a determin... algorithm, we define its competitive ratio ρ as

$$\rho = \max_{\sigma} \frac{B_{\mathrm{OPT}}(\sigma)}{B_A(\sigma)},$$

where the maximum is taken over all possible sequences of calls. If A is a randomized algorithm, we d... its competitive ratio ρ as

$$\rho = \max_{\sigma} \frac{B_{\mathrm{OPT}}(\sigma)}{\mathcal{E}[B_A(\sigma)]},$$

where $\mathcal{E}[B_A(\sigma)]$ is the expectation of the number of calls accepted by A, and the maximum is taken... all possible sequences of calls.

We compare the performance of deterministic algorithms against *offline adversaries*, that is, adv... that have knowledge of the behavior of the deterministic algorithm in advance. In the case of rando... algorithms, we consider *oblivious adversaries* whose knowledge is limited to the probability distri... of the random choices of the randomized algorithm.

In the next sections, we survey related work on the call control problem. We adapt to wireless cellular networks ideas that have been proposed in similar context (such as call control in optical networks [1,2,22]) and present recent work of the authors [5–8]. Presenting very briefly the related results, deterministic algorithms do not provide efficient online solutions to the problem. The greedy algorithm (studied in References 19 and 20) is probably the simplest such algorithm. Simple arguments can show that its competitive ratio is at least 3 in cellular networks and, furthermore, no deterministic algorithm has better competitive ratio. So, better solutions are possible only through randomized algorithms. Adapting the classify and randomly select idea from Reference 2, a simple 3-competitive algorithm can be achieved. Although not beating the lower bound for deterministic algorithms, unlike the greedy algorithm, this algorithm achieves competitiveness of 3 even in networks with arbitrarily many frequencies.

Improving the bound of 3 was the aim of previous work of the authors. As a first step, Reference 6 studies the randomized algorithm p-RANDOM, which marginally deviates from the behavior of the greedy algorithm. Intuitively, p-RANDOM accepts a call with probability p whenever this is possible. Although simple, this idea is powerful enough to cross the barrier of 3 in single-frequency cellular networks. The analysis is interesting, as well. In order to account for the expected benefit of the algorithm, the benefit is amortized to optimal calls and lower bounds on the expectation of the amortized benefit per optimal call are enough to prove upper bounds on the competitive ratio. For obtaining good bounds, detailed case analysis is required.

In another step, in order to investigate the power of randomization, the authors in References 7 and 8 develop simple randomized algorithms extending the "classify and randomly select" paradigm. These algorithms are based on colorings of the interference graph satisfying certain properties. Besides improving further the upper bounds on the competitiveness of randomized call control, these algorithms are particularly simple since they use only a small constant number of random bits or comparably weak random sources. On the other hand, p-RANDOM requires as many bits as the number of calls in a sequence. Furthermore, the "classify and randomly select" based algorithms work for cellular networks supporting many frequencies without any overhead in performance.

In the rest of the chapter, we survey the above results. We first present in Section 38.2 the greedy algorithm and the "classify and randomly select" algorithm produced by naively applying the ideas of Reference 2 in a cellular network and present the simple argument that proves the lower bound on the competitive ratio of deterministic algorithms. As part of the analysis of the greedy algorithm, we present an interesting technique from Reference 1 for transforming call control algorithms for one frequency to algorithms for many frequencies with a small sacrifice on performance. In Section 38.3, we present the amortized benefit analysis of algorithm p-RANDOM by presenting the related case analysis. The "classify and randomly select"-based algorithms are presented in Section 38.4. In Section 38.5, we demonstrate how Yao's principle can be used in order to prove lower bounds for randomized algorithms against oblivious adversaries. We conclude by briefly discussing interesting extensions of our model and open problems.

38.2 Two Simple Algorithms

In this section, we describe two well-known online algorithms for call control in wireless networks: the greedy algorithm and a randomized algorithm based on the "classify and randomly select" paradigm. Also, we present a lower bound on the competitiveness of deterministic online call control algorithms. These results will be the starting point for the improvements we will present in the next sections.

Assume that a sequence of calls σ appears in a network that support w frequencies $1, 2, \ldots, w$. The greedy algorithm is an intuitive deterministic algorithm. For any new call c at a cell v, the greedy algorithm searches for the minimum available frequency, that is, for the minimum frequency among frequencies $1, 2, \ldots, w$ that has not been assigned to calls in cell v or its adjacent cells. If such a frequency exists, the call c is accepted and is assigned this frequency; otherwise, the call is rejected.

Pantziou et al. [20] have proved that this algorithm is at most $(\Delta + 1)$-competitive against offline adversaries for networks supporting many frequency, where Δ is the degree of the network. The following statement slightly extends this result.

Theorem 38.1 *Let $G = (V, E)$ be an interference graph, v a vertex of G, and Γ_v the maximum independent set in the neighborhood of v. The greedy algorithm is $\frac{1}{1-\exp^{-1/\gamma}}$-competitive against an offline adversary, where $\gamma = \max_{v \in V} |\Gamma_v|$.*

To prove this statement, we will first show that the greedy algorithm is γ-competitive if the network supports only one frequency. Given a sequence of calls, denote by B_A be the set of calls accepted by the greedy algorithm and B_{OPT} the set of calls accepted by the optimal algorithm. Observe that for each optimal call c not accepted by the algorithm, there are at most γ calls in $B_{\mathrm{OPT}} \backslash B_A$ that are rejected because of the acceptance of c. Thus,

$$\begin{aligned} |B_{\mathrm{OPT}}| &= |B_{\mathrm{OPT}} \backslash B_A| + |B_{\mathrm{OPT}} \cap B_A| \\ &\leq \gamma |B_A \backslash B_{\mathrm{OPT}}| + |B_{\mathrm{OPT}} \cap B_A| \\ &\leq \gamma |B_A|. \end{aligned}$$

To prove Theorem 38.1, we will use a technique of Awerbuch et al. [1] who present a simple way for transforming call control algorithms designed for networks with one frequency to call control algorithms for networks with arbitrarily many frequencies, with a small sacrifice in competitiveness. Consider a wireless cellular network and a (deterministic or randomized) online call control algorithm ALG-1 for one frequency. A call control algorithm ALG for w frequencies can be constructed in the following way. For each call c, we execute the algorithm ALG-1 for each of the w frequencies until either c is accepted or the frequency spectrum is exhausted (and the call c is rejected), that is,

```
for any new call c
    for i = 1 to w
        run ALG-1(c) for frequency i
        if c was accepted then
            assign frequency i to c
            stop
    reject c.
```

Lemma 38.1 *(Awerbuch et al. [1]) If ALG-1 is ρ-competitive, then ALG has competitive ratio*

$$\frac{1}{1 - \left(1 - \frac{1}{\rho w}\right)^{w}} \leq \frac{1}{1 - \exp(-1/\rho)}.$$

Hence, Theorem 38.1 immediately follows by Lemma 38.1 and the fact that the greedy algorithm is γ-competitive for one frequency. For cellular networks, where the interference graph is a hexagon graph, it is $\gamma = 3$, and Theorem 38.1 yields the following corollary.

Corollary 38.1 *The greedy algorithm is 3.53-competitive in cellular networks with arbitrarily many frequencies.*

Note that γ is a lower bound for the competitive ratio of every deterministic algorithm. Consider a network that supports one frequency and consists of a cell v and γ mutually non-adjacent cells $v_1, v_2, \ldots, v_\gamma$ that are adjacent to v. Consider, now, the following sequence of calls produced by an adversary that has knowledge of the way that the algorithm makes its decisions. First, a call c is presented in cell v. If the algorithm rejects c, then the adversary stops the sequence. In this case, the algorithm has no benefit from its execution. If the algorithm accepts the call c, the adversary presents γ calls $c_1, c_2, \ldots, c_\gamma$ in cells $v_1, v_2, \ldots, v_\gamma$, respectively. The benefit of the algorithm is then 1 while the optimal algorithm would obtain benefit γ by rejecting call c and accepting the calls $c_1, \ldots, c_\gamma$. Adapting this argument to cellular networks, we obtain the following statement.

Theorem 38.2 *No deterministic algorithm can be better than 3-competitive against an offline adversary.*

Obviously, the best the algorithm A can do is to accept all calls presented in cells that are nonadjacent to cells where previously accepted calls are located. But this is exactly what the greedy algorithm does for networks that support one frequency.

The "classify and randomly select" paradigm (introduced in a different context in Reference 2; see also Reference 20) uses a coloring of the cells of the network (coloring of the interference graph) with positive integer (colors) $1, 2, \ldots$ in such way that adjacent cells are assigned different colors. The randomized algorithm classifies the calls of the sequence into a number of classes; class i contains calls appeared in cells colored with color i. It then selects uniformly at random one of the classes, and considers only calls that belong to the selected class, rejecting all other calls. Once a call of the selected class appears, the greedy algorithm is used.

Using simple arguments, we can prove that the "classify and randomly select" algorithm CRS is χ-competitive against oblivious adversaries, where χ is the number of colors used in the coloring of the cells of the network. This may lead to algorithms with competitive ratio equal to the chromatic number (and no better, in general) of the corresponding interference graph, given that an optimal coloring (i.e., with the minimum number of colors) is available. In cellular networks, the interference graph is 3-colorable. Hence, we obtain the following statement.

Theorem 38.3 *Algorithm CRS is 3-competitive against oblivious adversaries in cellular networks supporting arbitrarily many frequencies.*

38.3 The Algorithm p-Random

As a first attempt in order to prove that randomization indeed helps in order to beat the lower bound of 3 on the competitiveness of deterministic algorithms, we present and analyze the algorithm p-Random, a randomized call control algorithm for cellular networks that supports one frequency. Algorithm p-Random receives as input a sequence of calls in an online manner and works as follows:

```
1. Initially, all cells are unmarked.
2. for any new call c in a cell v
3.     if v is marked then reject c.
4.     if v has an accepted call or is adjacent to a cell
       with an accepted call, then reject c
5.     else
6.         with probability p accept c.
7.         with probability 1 − p reject c and mark v.
```

The algorithm uses a parameter $p \in [1/3, 1]$. Obviously, if it is $p < 1/3$, the competitive ratio will be greater than 3, since the expected benefit of the algorithm on a sequence of a single call will be p. The algorithm is simple and can be easily implemented with small communication overhead (exchange of messages) between the base stations of the network.

Marking cells on rejection guarantees that algorithm p-Random does not simulate the greedy deterministic one. Assume otherwise, that marking is not used. Then, consider an adversary that presents t calls in a cell v and one call in 3 (mutually not adjacent) cells adjacent to v. The probability that the randomized algorithm does not accept a call in cell v drops exponentially as t increases, and the benefit approaches 1, while the optimal benefit is 3.

Note that algorithm p-Random may accept at most one call in each cell, but this is also the case for any algorithm running in networks that support one frequency (including the optimal one). Thus, for the competitive analysis of algorithm p-Random, we will only consider sequences of calls with at most one call per cell. Also, there is no need for taking into account the procedure of marking cells during the analysis.

We now prove the upper bound on the competitive ratio of algorithm p-Random as a function of p. Our main statement is the following.

Theorem 38.4 *For $p \in [1/3, 1]$, algorithm p-Random has competitive ratio at most*

$$\frac{3}{5p - 7p^2 + 3p^3}$$

against oblivious adversaries.

Proof. Let σ be a sequence of calls. We assume that σ has been fixed in advance and will be revealed to the algorithm in an online manner. We make this assumption because we are interested in the competitiveness of the algorithm against oblivious adversaries whose knowledge is limited to the probability distribution of the random choices of the algorithm (i.e., the parameter p).

Consider the execution of algorithm p-Random on σ. For any call $c \in \sigma$, we denote by $X(c)$ the random variable that indicates whether the algorithm accepted c. Clearly, the benefit of algorithm p-Random on σ can be expressed as

$$B(\sigma) = \sum_{c \in \sigma} X(c).$$

Let $A(\sigma)$ be the set of calls in σ accepted by the optimal algorithm. For each call $c \in A(\sigma)$, we define the amortized benefit $\bar{b}(c)$ as

$$\bar{b}(c) = X(c) + \sum_{c' \in \gamma(c)} \frac{X(c')}{d(c')},$$

where $\gamma(c)$ denotes the set of calls of the sequence in cells adjacent to c. For each call $c' \notin A(\sigma)$, $d(c')$ is the number of calls in $A(\sigma)$ that are in cells adjacent to the cell of c. By the two equalities above, it is clear that

$$B(\sigma) = \sum_{c \in A(\sigma)} \bar{b}(c).$$

Furthermore, note that for any call $c' \notin A(\sigma)$, $d(c') \leq 3$. We obtain that

$$\bar{b}(c) \geq X(c) + \frac{\sum_{c' \in \gamma(c)} X(c')}{3}$$

and, by linearity of expectation,

$$\mathcal{E}[B(\sigma)] \geq \sum_{c \in A(\sigma)} \left(\mathcal{E}\left[X(c) + \frac{\sum_{c' \in \gamma(c)} X(c')}{3} \right] \right) \tag{38.1}$$

Let $\gamma'(c)$ be the set of calls in cells adjacent to the cell of c that appear prior to c in the sequence σ. Clearly, $\gamma'(c) \subseteq \gamma(c)$, which implies that

$$\sum_{c' \in \gamma(c)} X(c') \geq \sum_{c' \in \gamma'(c)} X(c').$$

Thus, Equation 38.1 yields

$$\mathcal{E}[B(\sigma)] \geq \sum_{c \in A(\sigma)} \left(\mathcal{E}\left[X(c) + \frac{\sum_{c' \in \gamma'(c)} X(c')}{3} \right] \right) \tag{38.2}$$

In what follows, we will try to bound from below the expectation of the random variable $X(c) + \frac{\sum_{c' \in \gamma'(c)} X(c')}{3}$, for each call $c \in A(\sigma)$.

We concentrate on a call $c \in A(\sigma)$. Let $\Omega = 2^{\gamma'(c)}$ be the set that contains all possible subsets of $\gamma'(c)$. We define the *effective neighborhood* of c, denoted by $\Gamma(c)$, to be the subset of $\gamma'(c)$ that contains the calls of $\gamma'(c)$ which, when it appears, is unconstrained by calls of σ at distance 2 from c. Clearly, $\Gamma(c)$ is a random variable taking its values from the sample space Ω. Intuitively, whether an optimal call c is accepted by the algorithm depends on its effective neighborhood $\Gamma(c)$. We have

$$\begin{aligned}
&\mathcal{E}\left[X(c) + \frac{\sum_{c' \in \gamma'(c)} X(c')}{3} \right] = \\
&\sum_{\gamma \in \Omega} \mathcal{E}\left[X(c) + \frac{\sum_{c' \in \gamma} X(c')}{3} | \Gamma(c) = \gamma \right] \cdot \Pr[\Gamma(c) = \gamma] \geq \\
&\min_{\gamma \in \Omega} \left\{ \mathcal{E}\left[X(c) + \frac{\sum_{c' \in \gamma} X(c')}{3} | \Gamma(c) = \gamma \right] \right\} = \\
&\min_{\gamma \in \Omega} \left\{ \mathcal{E}[X(c) | \Gamma(c) = \gamma] + \frac{\mathcal{E}\left[\sum_{c' \in \gamma} X(c') | \Gamma(c) = \gamma \right]}{3} \right\}.
\end{aligned} \tag{38.3}$$

To compute $\mathcal{E}[X(c)|\Gamma(c) = \gamma]$, we observe that algorithm p-Random may accept c only if it has rejected all calls in its effective neighborhood γ. The probability that all calls of γ are rejected given that $\Gamma(c) = \gamma$ is $(1-p)^{|\gamma|}$, and then c is accepted with probability p. Thus,

$$\mathcal{E}[X(c)|\Gamma(c) = \gamma] = p(1-p)^{|\gamma|}. \tag{38.4}$$

We now bound from below $\mathcal{E}\left[\sum_{c' \in \gamma} X(c') | \Gamma(c) = \gamma\right]$ by distinguishing between cases according to the size of the effective neighborhood $|\gamma|$.

Claim 38.5 For all $p \in [1/3, 1]$,

$$\mathcal{E}\left[\sum_{c' \in \gamma} X(c') | \Gamma(c) = \gamma \right] \geq \begin{cases} 0 & \textit{if } |\gamma| = 0 \\ p & \textit{if } |\gamma| = 1 \\ 2p - p^2 & \textit{if } |\gamma| = 2 \\ 3p - 2p^2 & \textit{if } |\gamma| = 3 \\ 4p - 3p^2 + p^3 & \textit{if } |\gamma| = 4 \\ 5p - 4p^2 + p^3 & \textit{if } |\gamma| = 5 \\ 6p - 5p^2 + p^3 & \textit{if } |\gamma| = 6 \end{cases}$$

Proof. In Figures 38.2 through 38.6, we give all possible cases for the effective neighborhood of an optimal call c in a sequence of calls σ. In each figure the optimal call is denoted by the black circle in the middle cell while black circles in the outer cells denote calls in the effective neighborhood γ of c. An arrow from a call c_1 to another call c_2 indicates that c_1 appears in σ prior to c_2. In the figures, we have eliminated the symmetric cases.

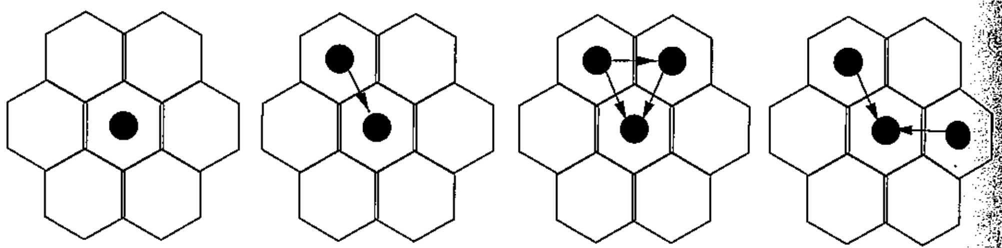

FIGURE 38.2 The cases $|\gamma| = 0, 1, 2$.

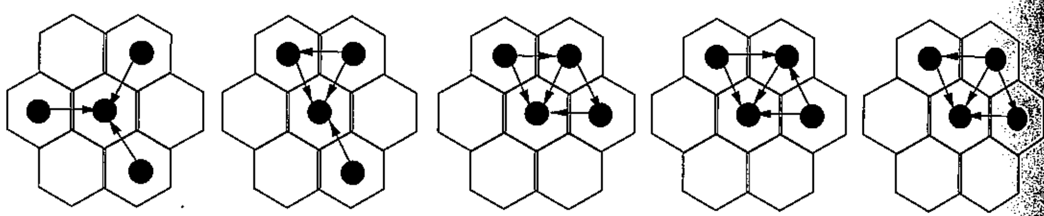

FIGURE 38.3 The case $|\gamma| = 3$.

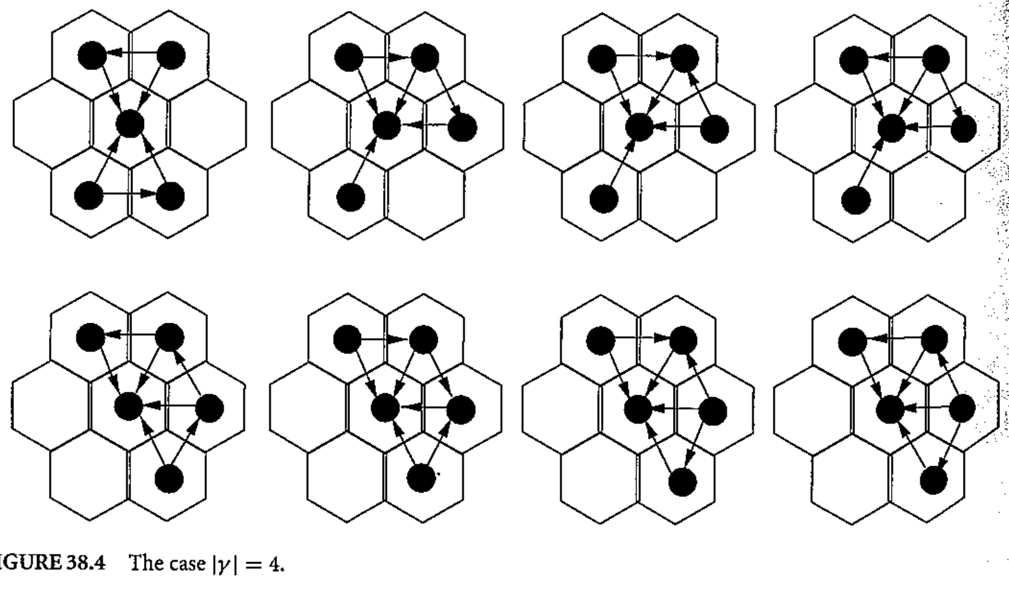

FIGURE 38.4 The case $|\gamma| = 4$.

The proof is trivial for the cases $|\gamma| = 0$ and $|\gamma| = 1$ (the two leftmost cases in Figure 38.2). In the third case of Figure 38.2 (where $|\gamma| = 2$), we observe that the algorithm accepts the first call in γ with probability p and the second one with probability $p(1-p)$. In total, the expectation of the number of accepted calls in γ is $2p - p^2$. In the rightmost case of Figure 38.2, the expectation of the number of accepted calls in γ is $2p < 2p - p^2$.

Similarly, we can compute the desired lower bounds on $\mathcal{E}\left[\sum_{c' \in \gamma} X(c') | \Gamma(c) = \gamma\right]$ for the cases $|\gamma| = 3, 4, 5, 6$. ■

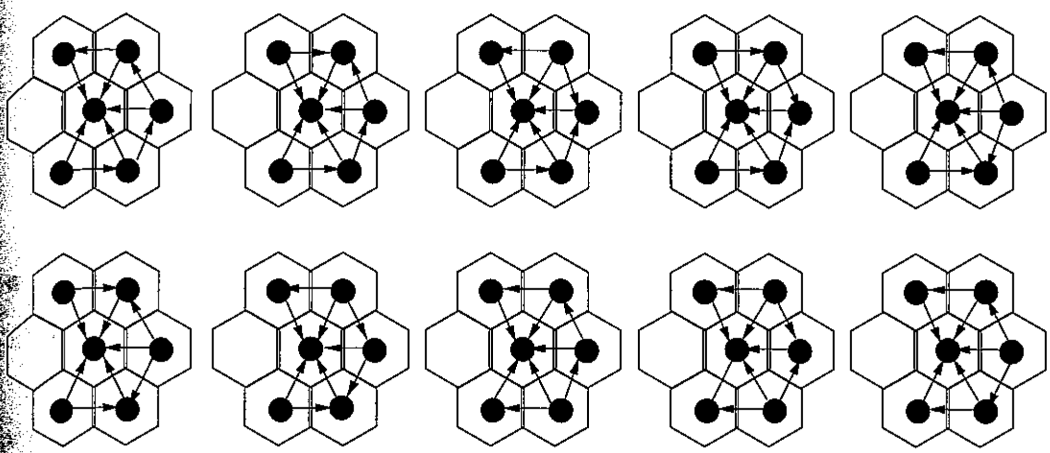

FIGURE 38.5 The case $|\gamma| = 5$.

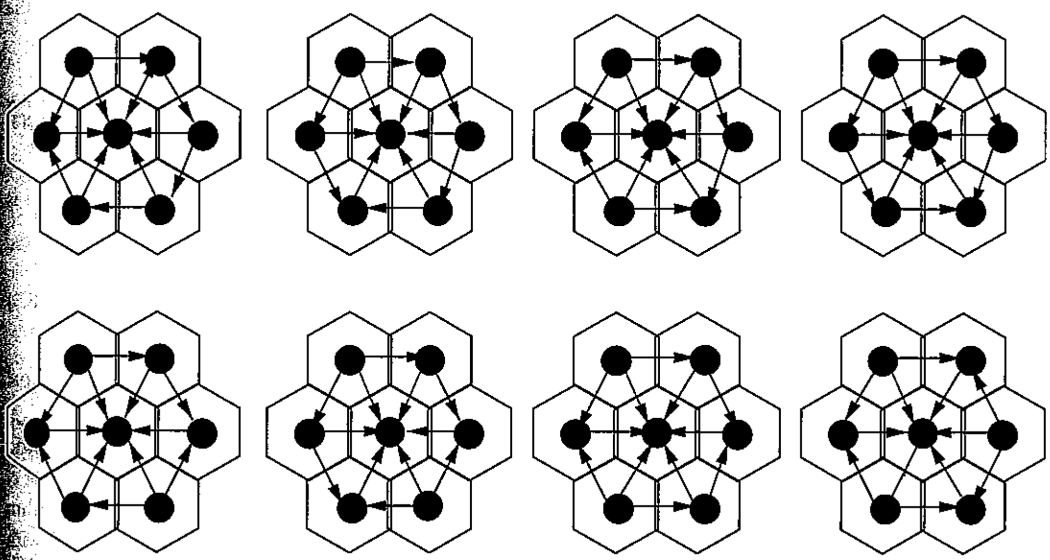

FIGURE 38.6 The case $|\gamma| = 6$.

By making calculations with Equations 38.3, 38.4, and Claim 38.5 , we obtain that

$$\mathcal{E}\left[X(c) + \frac{\sum_{c' \in \gamma'(c)} X(c')}{3}\right] \geq$$

$$\min_{\gamma \in \Omega}\left\{\mathcal{E}[X(c)|\Gamma(c) = \gamma] + \frac{\mathcal{E}\left[\sum_{c' \in \gamma} X(c')|\Gamma(c) = \gamma\right]}{3}\right\} \geq$$

$$\min_{\gamma \in \Omega : |\gamma| = 2}\left\{\mathcal{E}[X(c)|\Gamma(c) = \gamma] + \frac{\mathcal{E}\left[\sum_{c' \in \gamma} X(c')|\Gamma(c) = \gamma\right]}{3}\right\} \geq$$

$$p(1-p)^2 + \frac{2p - p^2}{3} = \frac{5p - 7p^2 + 3p^3}{3}.$$

Now, using Equation 38.2 we obtain that

$$\mathcal{E}[B(\sigma)] \geq \sum_{c \in A(\sigma)} \frac{5p - 7p^2 + 3p^3}{3}$$
$$= \frac{5p - 7p^2 + 3p^3}{3} \cdot B_{\mathrm{OPT}}(\sigma).$$

This completes the proof of Theorem 38.4. ∎

The expression in Theorem 38.4 is minimized to $729/265 = 2.651$ for $p = 5/9$. Thus, we obtain the following result.

Corollary 38.2 *There exists an online randomized call control algorithm for cellular networks with one frequency that is at most* 2.651*-competitive against oblivious adversaries.*

In the following, we show that our analysis is not far from being tight. In particular, we prove the following.

Theorem 38.6 *For any* $p \in (1/3, 1)$, *algorithm* p-RANDOM *is at least* 2.469*-competitive against oblivious adversaries.*

Proof. We will show that the competitive ratio of algorithm p-RANDOM against oblivious adversaries is at least

$$\max\left\{\frac{3}{4p - 3p^2}, \frac{3}{5p - 7p^2 + 4p^3 - p^4}\right\} \tag{38.5}$$

by constructing two sequences σ_1 and σ_2 of calls for which the competitive ratio of algorithm p-RANDOM is $\frac{3}{4p-3p^2}$ and $\frac{3}{5p-7p^2+4p^3-p^4}$, respectively.

Sequence σ_1 is depicted in the left part of Figure 38.7. In round 1, a call appears at some cell c, and in round 2 one call appears in each one of the three mutually adjacent cells in the neighborhood of c. Clearly, the benefit of the optimal algorithm is 3. To compute the expectation of the benefit of algorithm p-RANDOM, we observe that with probability p, the call presented in round 1 is accepted, and with probability $1 - p$, the call presented in round 1 is rejected and each of the three calls presented in round 2 is accepted with probability p. Thus, the expectation of the benefit of the algorithm on sequence σ_1 is $p + (1 - p)3p = 4p - 3p^2$.

Sequence σ_2 is depicted in the right part of Figure 38.7. Calls appear in four rounds. The labels on the calls denote the round in which the calls appear. Clearly, the benefit of the optimal algorithm is 18 since the optimal algorithm would accept the calls that appear in rounds 3 and 4. To compute the expectation of the benefit of algorithm p-RANDOM on sequence σ_2, we first compute the probability that each call is accepted.

- A call that appears in round 1 is accepted with probability p.
- A call that appears in round 2 can be accepted if its adjacent call that appeared in round 1 has been rejected; thus, the probability that a call that appears in round 2 is accepted is $p(1 - p)$.
- A call that appears in round 3 can be accepted if its adjacent calls that appeared in rounds 1 and 2 have been rejected; thus, the probability that a call that appears in round 3 is accepted is $p(1 - p)^2$.
- A call that appears in round 4 can be accepted if its adjacent calls that appeared in rounds 1 and 2 have been rejected. The probability that a call that appears in round 1 is rejected is $1 - p$ while the probability that a call that appears in round 2 is rejected is $1 - p - (1 - p)^2$. Thus, the probability that a call that appears in round 4 is accepted is $p(1 - p)\left(1 - p - (1 - p)^2\right)$.

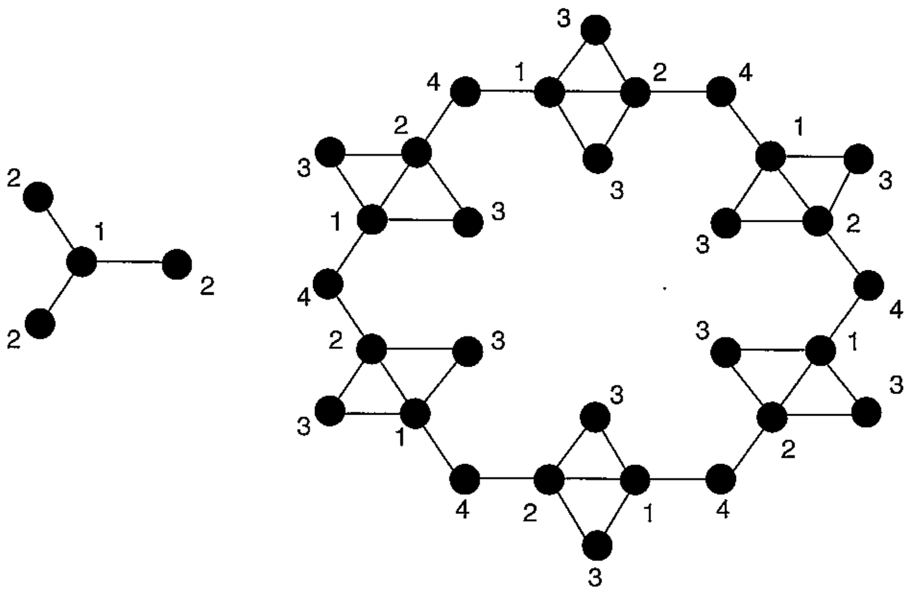

FIGURE 38.7 The lower bound on the performance of algorithm p-RANDOM.

Note the number of calls that appear in rounds 1, 2, 3, and 4 is 6, 6, 12, and 6, respectively. Thus, we obtain that the expectation of the benefit of the algorithm is

$$6p + 6p(1-p) + 12p(1-p)^2 + 6p(1-p)\left(1 - p - (1-p)^2\right) = 30p - 42p^2 + 24p^3 - 6p^4.$$

By the above constructions, it is clear that the competitive ratio of algorithm p-RANDOM is lower bounded by Equation 38.5. This expression is minimized for $p \approx 0.6145$ to 2.469. This completes the proof of the theorem. ■

Using algorithm p-RANDOM, a call control algorithm for w frequencies can be constructed using the technique of Awerbuch et al. [1] presented in Section 38.2. For each call c, we execute the algorithm p-RANDOM for each of the w frequencies until either c is accepted or the frequency spectrum is exhausted (and the call c is rejected). Using Lemma 38.1, we obtain that the competitive ratio we achieve in this way for two frequencies is 2.927. Unfortunately, for larger numbers of frequencies, the competitive ratio becomes larger than 3.

38.4 "Classify and Randomly Select"-Based Algorithms

In this section, we present "classify and randomly select"-based algorithms originally presented in References 7 and 8. We start with algorithm CRS-A that works in networks with one frequency and achieves a competitive ratio against oblivious adversaries similar (but slightly inferior) to that proved for algorithm p-RANDOM.

Algorithm CRS-A uses a coloring of the cells with four colors 0, 1, 2, and 3, such that only two colors are used in the cells belonging to the same axis. This can be done by coloring the cells in the same x-row with either the colors 0 and 1 or the colors 2 and 3, coloring the cells in the same y-row with either the colors 0 and 2 or the colors 1 and 3, and coloring the cells in the same z-row with either the colors 0 and 3 or the colors 1 and 2. Such a coloring is depicted in the left part of Figure 38.8.

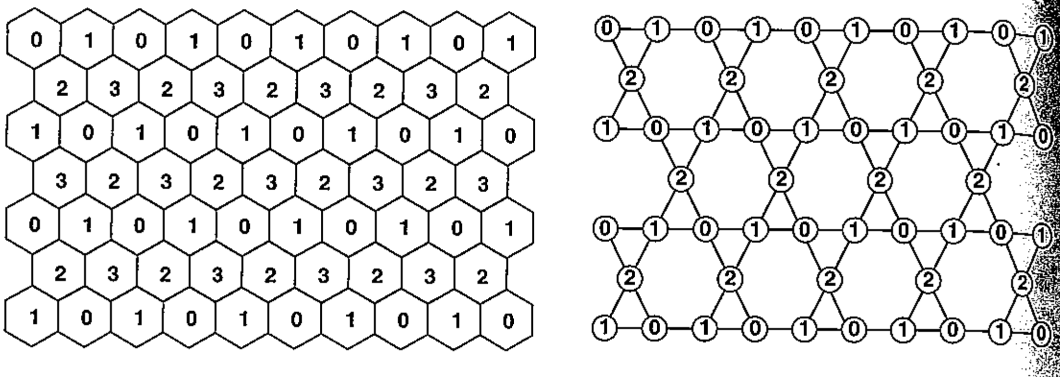

FIGURE 38.8 The 4-coloring used by algorithm CRS-A and the corresponding subgraph of the interference graph induced by the nodes not colored with color 3.

Algorithm CRS-A randomly selects one out of the four colors and executes the greedy algorithm on the cells colored with the other three colors, ignoring (i.e., rejecting) all calls in cells colored with the selected color.

Theorem 38.7 *Algorithm CRS-A in cellular networks supporting one frequency is* 8/3*-competitive against oblivious adversaries.*

Proof. Let σ be a sequence of calls and denote by O the set of calls accepted by the optimal algorithm. Denote by σ' the set of calls in cells that are not colored with the color selected and by O' the set of calls the optimal algorithm would have accepted on input σ'. Clearly, $|O'|$ will be at least as large as the subset of O that belongs to σ'. Since the probability that the cell of a call in O is not colored with the color selected is 3/4, it is $\mathcal{E}[|O'|] \geq 3/4|O|$.

Now let B be the set of calls accepted by algorithm CRS-A, that is, the set of calls accepted by the greedy algorithm when executed on sequence σ'. Observe that each call in O' either belongs in B or it is rejected because some other call is accepted. Furthermore, a call in $B \setminus O'$ can cause the rejection of at most two calls of O'. This implies that $|B| \geq |O'|/2$, which yields that the competitive ratio of algorithm CRS-A is

$$\frac{|O|}{\mathcal{E}[|B|]} \leq \frac{2|O|}{\mathcal{E}[|O'|]} \leq \frac{8}{3}.$$

∎

The main advantage of algorithm CRS-A is that it uses only two random bits. In the next section, we present simple online algorithms with improved competitive ratios that use slightly stronger random sources and work on networks with arbitrarily many frequencies.

Algorithm CRS-A can be seen as an algorithm based on the "classify and randomly select" paradigm. It uses a coloring of the interference graph (not necessarily using the minimum possible number of colors) and a classification of the colors. It starts by randomly selecting a color class (i.e., a set of colors) and then run the greedy algorithm in the cells colored with colors from this color class, ignoring (i.e., rejecting) calls in cells colored with colors not belonging to this class. Algorithm CRS-A uses a coloring of the interference graph with four colors 0, 1, 2, and 3, and the four color classes $\{0, 1, 2\}$, $\{0, 1, 3\}$, $\{0, 2, 3\}$, and $\{1, 2, 3\}$. Note that, in the previously known algorithms based on the "classify and randomly select" paradigm, color classes are singletons (e.g., References 1 and 20).

The following simple lemma gives a sufficient condition for obtaining efficient online algorithms based on the "classify and randomly select" paradigm.

Lemma 38.2 *Consider a network with interference graph $G = (V, E)$ that supports w frequencies and let χ be a coloring of the nodes of V with the colors of a set X. If there exist ν sets of colors $s_0, s_1, \ldots, s_{\nu-1} \subseteq X$ and an integer λ such that*

- *Each color of X belongs to at least λ different sets of the sets $s_0, s_1, \ldots, s_{\nu-1}$.*
- *For $i = 0, 1, \ldots, \nu - 1$, each connected component of the subgraph of G induced by the nodes colored with colors in s_i is a clique.*

then there exists an online randomized call control algorithm for the network G, which has competitive ratio ν/λ against oblivious adversaries.

Proof. Consider a network with interference graph G that supports w frequencies and the randomized online algorithm working as follows. The algorithm randomly selects one out of the ν color classes $s_0, \ldots, s_{\nu-1}$ and executes the greedy algorithm on the cells colored with colors of the selected class, rejecting all calls in cells colored with colors not in the selected class.

Let σ be a sequence of calls and let O be the set of calls accepted by the optimal algorithm on input σ. Assume that the algorithm selects the color class s_i. Let σ' be the sequence of calls in cells colored with colors in s_i, and O' be the set of calls accepted by the optimal algorithm on input σ'. Also, we denote by B the set of calls accepted by the algorithm.

First, we can easily show that $|B| = |O'|$. Let G_j be a connected component of the subgraph of G induced by the nodes of G colored with colors in s_i. Let σ_j be the subsequence of σ' in cells corresponding to nodes of G_j. Clearly, any algorithm (including the optimal one) will accept at most one call of σ_j at each frequency. If the optimal algorithm accepts w calls, this means that the sequence σ_j has at least w calls and the greedy algorithm, when executed on σ', will accept w calls from σ_j (one call in each one of the available frequencies). If the optimal algorithm accepts $w' < w$ calls from σ_j, this means that σ_j contains exactly $w' < w$ calls and the greedy algorithm will accept them all in w' different frequencies. Since a call of σ_j is not constrained by a call in $\sigma_{j'}$ for $j \neq j'$, we obtain that $|B| = |O'|$.

The proof is completed by observing that the expected benefit of the optimal algorithm on input σ' over all possible sequences σ' defined by the random selection of the algorithm is $\mathcal{E}[|O'|] \geq \frac{\nu}{\lambda}|O|$, since, for each call in O, the probability that the color of its cell belongs to the color class selected is at least ν/λ. Hence, the competitive ratio of the algorithm against oblivious adversaries is

$$\frac{|O|}{\mathcal{E}[|B|]} = \frac{|O|}{\mathcal{E}[|O'|]} \leq \lambda/\nu.$$

■

Next, we present simple randomized online algorithms for call control in cellular networks, namely, CRS-B, CRS-C, and CRS-D, which are also based on the "classify and randomly select" paradigm and achieve even better competitive ratios.

Consider a coloring of the cells with five colors 0, 1, 2, 3, and 4 such that for each $i \in \{0, 1, 2, 3, 4\}$, and for each cell colored with color i, the two adjacent cells in the same x-row are colored with colors $(i-1) \bmod 5$ and $(i+1) \bmod 5$, while the remaining four of its adjacent cells are colored with colors $(i+2) \bmod 5$ and $(i+3) \bmod 5$. Such a coloring is depicted in the left part of Figure 38.9. Also, define $s_i = \{i, (i+1) \bmod 5\}$, for $i = 0, 1, \ldots, 4$. Observe that, for each $i = 0, 1, \ldots, 4$, each pair of adjacent cells colored with the colors i and $(i+1) \bmod 5$ is adjacent to cells colored with colors $(i+2) \bmod 5$, $(i+3) \bmod 5$, and $(i+4) \bmod 5$, that is, colors not belonging to s_i. Thus, the coloring together with the color classes s_i satisfy the conditions of Lemma 38.2 with $\nu = 5$ and $\lambda = 2$. We call CRS-B the algorithm that uses this coloring and works according to the "classify and randomly select" paradigm as in the proof of Lemma 38.2. We obtain the following:

Theorem 38.8 *Algorithm CRS-B in cellular networks is 5/2-competitive against oblivious adversaries.*

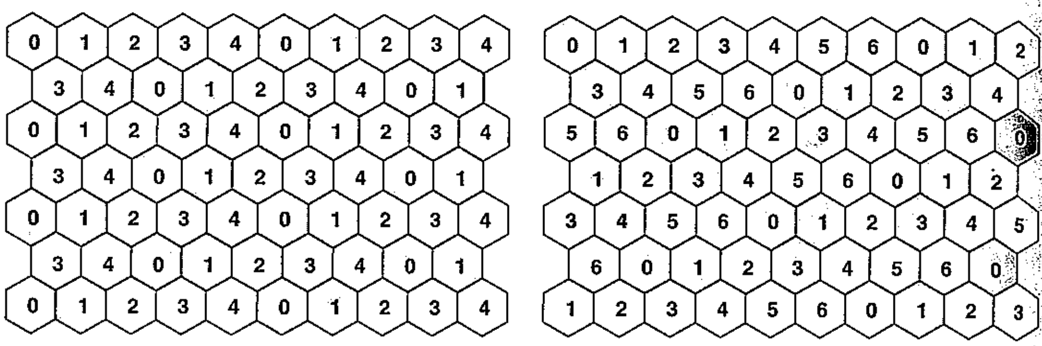

FIGURE 38.9 The 5-coloring used by algorithm CRS-B and the 7-coloring used by algorithm CRS-C. The gray cells are those colored with the colors in set s_0.

Now consider a coloring of the cells with seven colors $0, 1, \ldots, 6$ such that for each cell colored with color i (for $i = 0, \ldots, 6$), its two adjacent cells in the same x-row are colored with the colors $(i-1) \bmod 7$ and $(i+1) \bmod 7$, while its two adjacent cells in the same z-row are colored with colors $(i-3) \bmod 7$ and $(i+3) \bmod 7$. Such a coloring is depicted in the right part of Figure 38.9. Also, define $s_i = \{i, (i+1) \bmod 7, (i+3) \bmod 7\}$, for $i = 0, 1, \ldots, 6$. Observe that, for each $i = 0, 1, \ldots, 6$, each triangle of cells colored with the colors i, $(i+1) \bmod 7$, and $(i+3) \bmod 7$ is adjacent to cells colored with colors $(i+2) \bmod 7$, $(i+4) \bmod 7$, $(i+5) \bmod 7$, and $(i+6) \bmod 7$, that is, colors not belonging to s_i. Thus, the coloring together with the color classes s_i satisfy the conditions of Lemma 38.2 with $\nu = 7$ and $\lambda = 3$. We call CRS-C the algorithm that uses this coloring and works according to the "classify and randomly select" paradigm as in the proof of Lemma 38.2. We obtain the following.

Theorem 38.9 *Algorithm CRS-C in cellular networks is 7/3-competitive against oblivious adversaries.*

Algorithm CRS-D uses a coloring of the cells with 16 colors $0, \ldots, 15$ defined as follows. The cell with coordinates $(x, y, x+y)$ is colored with color $4(x \bmod 4) + y \bmod 4$. The color classes are defined as s_{4i+j} for $0 \le i, j \le 3$ as follows:

$$\begin{aligned} s_{4i+j} = \{&4i+j, 4i+(j+1) \bmod 4, 4((i+1) \bmod 4)+j,\\ &4((i+1) \bmod 4)+(j+2) \bmod 4, 4((i+2) \bmod 4)+(j+1) \bmod 4,\\ &4((i+2) \bmod 4)+(j+2) \bmod 4, 4((i+3) \bmod 4)+(j+3) \bmod 4\}. \end{aligned}$$

An example of this coloring is depicted in Figure 38.10.

We now show that the coloring and the color classes used by algorithm CRS-D satisfy the conditions of Lemma 38.2. Each color $k = 0, 1, \ldots, 15$ belongs to 7 of the 16 color classes $s_0, s_1, \ldots, s_{15}$; color $4i+j$ belongs to the color classes $4i+j$, $4(i+1)+(j+1)$, $4(i+2)+(j+2)$, $4(i+2)+(j+3)$, $4(i+3)+j$, $4(i+3)+(j+2)$, and $4i+(j+3) \bmod 4$ for $0 \le i, j \le 3$. Now, consider the cells colored with colors from the color class s_{4i+j} and the corresponding nodes of the interference graph. The connected components of the subgraph of the interference graph defined by these nodes are of the following types:

- Cliques of three nodes corresponding to cells colored with colors $4i+j$, $4i+(j+1) \bmod 4$, and $4((i+1) \bmod 4)+j$, respectively. Indeed, the neighborhood of such nodes contains nodes colored with colors $4i+(j+2) \bmod 4$, $4i+(j+3) \bmod 4$, $4(i+1) \bmod 4+(j+1) \bmod 4$, $4(i+1) \bmod 4+(j+3) \bmod 4$, $4(i+3) \bmod 4+j$, $4(i+3) \bmod 4+(j+1) \bmod 4$, $4(i+3) \bmod 4+(j+2) \bmod 4$, $4(i+2) \bmod 4+j$, and $4(i+2) \bmod 4+(j+3) \bmod 4$ that do not belong to class s_{4i+j}.
- Cliques of three nodes corresponding to cells colored with colors $4((i+1) \bmod 4)+(j+2) \bmod 4$, $4((i+2) \bmod 4)+(j+1) \bmod 4$, and $4((i+2) \bmod 4)+(j+2) \bmod 4$,

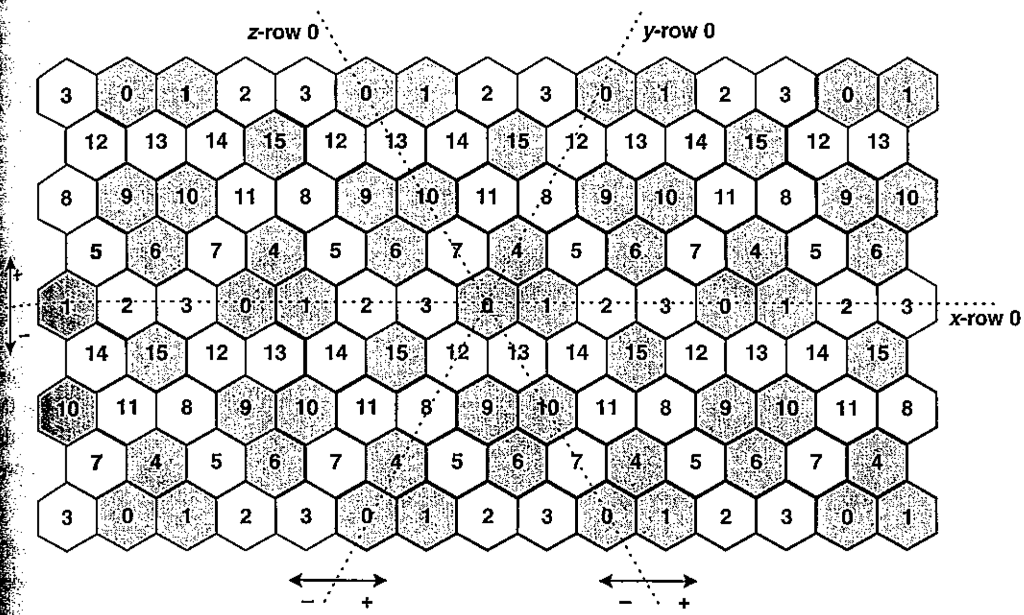

FIGURE 38.10 The 16-coloring used by algorithm CRS-D. The grey cells are those colored with colors in the class s_0.

respectively. Again, the neighborhood of such nodes contains nodes colored with colors $4i + (j + 2) \bmod 4$, $4i + (j + 3) \bmod 4$, $4(i + 1) \bmod 4 + (j + 1) \bmod 4$, $4(i + 1) \bmod 4 + (j+3) \bmod 4$, $4(i+3) \bmod 4+j$, $4(i+3) \bmod 4+(j+1) \bmod 4$, $4(i+3) \bmod 4+(j+2) \bmod 4$, $4(i + 2) \bmod 4 + j$, and $4(i + 2) \bmod 4 + (j + 3) \bmod 4$ that do not belong to class s_{4i+j}.

- Isolated nodes corresponding to cells colored with color $4((i + 3) \bmod 4) + (j + 3) \bmod 4$. The neighborhood of such a cell consists of cells colored with colors $4i+(j+2) \bmod 4$, $4i+(j+3) \bmod 4$, $4(i + 2) \bmod 4 + j$, $4(i + 2) \bmod 4 + (j + 3) \bmod 4$, $4(i + 3) \bmod 4 + j$, and $4(i + 3) \bmod 4 + (j + 2) \bmod 4$ that do not belong to class s_{4i+j}.

Hence, the coloring and the color classes used by algorithm CRS-D satisfy the conditions of Lemma 38.2 for $\lambda = 7$ and $\nu = 16$. This yields the following:

Theorem 38.10 *Algorithm CRS-D for call control in cellular networks is* 16/7*-competitive against oblivious adversaries.*

Obviously, the algorithm uses only 4 random bits for selecting equiprobably one out of the 16 color classes.

38.5 A Lower Bound

The randomized algorithms presented in the previous section significantly beat the lower bound on the competitiveness of deterministic algorithms. In what follows, using the Minimax Principle [23] (see also Reference 16), we prove a lower bound on the competitive ratio, against oblivious adversaries, of any randomized algorithm for call control in cellular networks. We consider networks that support one frequency; our lower bounds can be easily extended to networks that support multiple frequencies. In our proof, we use the following lemma.

Lemma 38.3 *(Minimax Principle, Yao [16]) Given a probability distribution $\mathcal{P}$ over sequences of calls σ, denote by $\mathcal{E}_{\mathcal{P}}[B_A(\sigma)]$ and $\mathcal{E}_{\mathcal{P}}[B_{\text{OPT}}(\sigma)]$ the expected benefit of a deterministic algorithm A and the optimal offline algorithm on sequences of calls generated according to $\mathcal{P}$. Define the competitiveness of A under $\mathcal{P}$, $c_A^{\mathcal{P}}$ to be such that*

$$c_A^{\mathcal{P}} = \frac{\mathcal{E}_{\mathcal{P}}[B_{\text{OPT}}(\sigma)]}{\mathcal{E}_{\mathcal{P}}[B_A(\sigma)]}.$$

Let A_R be a randomized algorithm. Then, the competitiveness of A under $\mathcal{P}$ is a lower bound on the competitive ratio of A_R against an oblivious adversary, that is, $c_A^{\mathcal{P}} \leq c_{A_R}$.

So, in order to prove a lower bound for any randomized algorithm, it suffices to define an adversary that produces sequences of calls according to a probability distribution and prove that the ratio of the expected optimal benefit over the expected benefit of any deterministic algorithm (which may know the probability distribution in advance) is above some value; by Lemma 38.3, this value will also be a lower bound for any randomized algorithm against oblivious adversaries.

Theorem 38.11 *No randomized online call control algorithm can be better than 2-competitive against oblivious adversaries in cellular networks.*

Proof. We present an adversary $\mathcal{ADV}$-2 that produces sequences of calls according to a probability distribution $\mathcal{P}_2$ that yields the lower bound. We show that the expected benefit of every deterministic algorithm (which may know $\mathcal{P}_2$ in advance) for such sequences of calls is at most 2, whereas the expected optimal benefit is at least 4. Then, Theorem 38.11 follows by Lemma 38.3. First, we describe the sequences of calls produced by $\mathcal{ADV}$-2 without explicitly giving the cells where they appear; then, we show how to construct them in a cellular network.

We start by defining a simpler adversary $\mathcal{ADV}$-1 that works as follows: It first produces two calls in two nonadjacent cells v_0 and v_1. Then it tosses a fair coin.

- On HEADS, it produces two calls in cells v_{00} and v_{01}, which are mutually nonadjacent, adjacent to v_0 and nonadjacent to v_1. Then, it stops.
- On TAILS, it produces two calls in cells v_{10} and v_{11}, which are mutually nonadjacent, adjacent to v_1 and nonadjacent to v_0. Then, it stops.

Now, consider the set of all possible deterministic algorithms $\mathcal{A}_1$ working on the sequences produced by $\mathcal{ADV}$-1. Such an algorithm $A_1 \in \mathcal{A}_1$ may follow one of the following strategies:

- It may accept both calls in cells v_0 and v_1, presented at the first step. This means that the calls presented in the second step cannot be accepted.
- It may reject both calls in cells v_0 and v_1, and then either accept one or both calls presented in the second step or reject them both.
- It may accept only one of the two calls in cells v_0 and v_1, and, if the calls produced at the second step by $\mathcal{ADV} - 1$ are nonadjacent to the accepted call, either accept one or both calls presented in the second step or reject them both.

In the first two cases, the expected benefit of the algorithm A_1 is at most 2. In the third case, the expected benefit is 1 (in the first step) plus the expected benefit in the second step. The latter is either zero with probability 1/2 (this is the case where the cells of the calls produced by the adversary in the second step are adjacent the cell of the call accepted by the algorithm in the first step) or at most 2 with probability 1/2. Overall, the expected benefit of the algorithm is at most 2.

The adversary $\mathcal{ADV}$-2 works as follows: It first produces two calls in two nonadjacent cells v_0 and v_1. Then it tosses a fair coin.

- On HEADS, it produces two calls in cells v_{00} and v_{01}, which are mutually nonadjacent, adjacent to v_0, and nonadjacent to v_1. Then, it tosses a fair coin.
  - On HEADS, it produces two calls in cells v_{000} and v_{001}, which are mutually nonadjacent, adjacent to v_0 and v_{00}, and nonadjacent to v_1 and v_{01}. Then, it stops.
  - On TAILS, it produces two calls in cells v_{010} and v_{011}, which are mutually nonadjacent, adjacent to v_0 and v_{01}, and nonadjacent to v_1 and v_{00}. Then, it stops.
- On TAILS, it produces two calls in cells v_{10} and v_{11}, which are mutually nonadjacent, adjacent to v_1, and nonadjacent to v_0. Then, it tosses a fair coin.
  - On HEADS, it produces two calls in cells v_{100} and v_{101}, which are mutually nonadjacent, adjacent to v_1 and v_{10}, and nonadjacent to v_0 and v_{11}. Then, it stops.
  - On TAILS, it produces two calls in cells v_{110} and v_{111}, which are mutually nonadjacent, adjacent to v_1 and v_{11}, and nonadjacent to v_0 and v_{10}. Then, it stops.

Observe that, the subsequence of the last four calls produced by $\mathcal{ADV}$-2 essentially belongs to the set of sequences of calls produced by $\mathcal{ADV}$-1.

Now, consider the set of all possible deterministic algorithms $\mathcal{A}_2$ working on the sequences produced by $\mathcal{ADV}$-2. Such an algorithm $A_2 \in \mathcal{A}_2$ may follow one of the following strategies:

- It may accept both calls in cells v_0 and v_1 presented at the first step. This means that the calls presented in the next steps cannot be accepted.
- It may reject both calls in cells v_0 and v_1 and then apply a deterministic algorithm A_1 on the subsequence presented after the first step.
- It may accept only one of the two calls in cells v_0 and v_1, and, then, if the calls produced at the next steps by $\mathcal{ADV}$-2 are nonadjacent to the accepted call, apply a deterministic algorithm A_1 on the subsequence presented after the first step.

In the first case, the expected benefit of the algorithm A_2 is at most 2. In the second case, the expected benefit of A_2 is the expected benefit of A_1 on the sequence of calls presented after the first step, that is, at most 2. In the third case, the expected benefit is 1 (in the first step) plus the expected benefit in the next steps. The benefit of the algorithm in the next steps is either zero with probability 1/2 (this is the case where the cells of the calls produced by the adversary in the next steps are nonadjacent to the cell of the call accepted by the algorithm in the first step) or the expected benefit of A_1 on the sequence of calls presented after the first step, that is, at most 2 with probability 1/2. Overall, the expected benefit of the algorithm is at most 2.

Furthermore, the expected optimal benefit on sequences produced by $\mathcal{ADV}$-2 is at least 4. Indeed, in each of the possible sequences

$$\sigma_2^{00} = \langle v_0, v_1, v_{00}, v_{01}, v_{000}, v_{001} \rangle, \quad \sigma_2^{01} = \langle v_0, v_1, v_{00}, v_{01}, v_{010}, v_{011} \rangle,$$
$$\sigma_2^{10} = \langle v_0, v_1, v_{10}, v_{11}, v_{100}, v_{101} \rangle, \quad \sigma_2^{11} = \langle v_0, v_1, v_{10}, v_{11}, v_{110}, v_{111} \rangle$$

generated by $\mathcal{ADV}$-2, the calls in cells $\langle v_1, v_{01}, v_{000}, v_{001} \rangle$, $\langle v_1, v_{00}, v_{010}, v_{011} \rangle$, $\langle v_0, v_{11}, v_{100}, v_{101} \rangle$, and $\langle v_0, v_{10}, v_{110}, v_{111} \rangle$ can be accepted, respectively. Overall, the ratio of the expected optimal benefit over the expected benefit of algorithm A_2 on the sequences generated by the adversary $\mathcal{ADV}$-2 is at least 2, which (by Lemma 38.3) is a lower bound on the competitive ratio of any randomized algorithm for call control.

Figure 38.11 shows how to locate the cells used by $\mathcal{ADV}$-2 so that all the above restrictions hold. This completes the proof of the theorem. ■

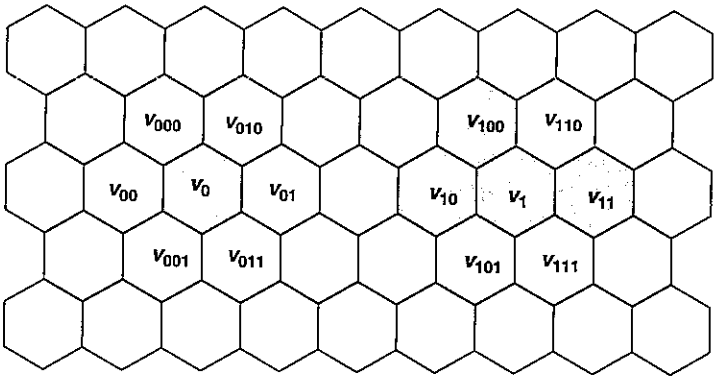

FIGURE 38.11 The calls that may be produced by the adversary $\mathcal{ADV}$-2. The gray cells host calls of the sequence σ_2^{10}.

38.6 Extensions and Open Problems

We conclude this chapter by briefly discussing extensions of our model and open problems. The cellular networks described so far have the property that the same frequency can be safely assigned to calls in cells at distance of at least 2 from each other so that no signal interference arises. These networks are said to be of *reuse distance* 2. We can generalize this constraint and consider cellular networks of reuse distance $k > 2$. Using counterexamples similar to the one described in Section 38.2, we can show that the competitive ratio of deterministic algorithms is at least 4 and 5 in cellular networks of reuse distance $k \in \{3, 4, 5\}$ and $k \geq 6$, respectively. So, investigating randomized online algorithms for these networks is interesting as well. The techniques of Section 38.4 can be used to beat the deterministic lower bounds. Note that Lemma 38.2 is general enough to be useful also in this case. In Reference 7, we present colorings of the interference graphs of these networks that satisfy the requirements of Lemma 38.2. These constructions can be thought of as generalizations of the coloring used by algorithm CRS-C presented in Section 38.4. The interested reader may refer to Reference 7 for a description of these colorings. Here, we only state the corresponding result.

Theorem 38.12 *There exists an online randomized call control algorithm that uses at most* $O(\log k)$ *random bits and has competitive ratio* $4 - \epsilon(k)$ *against oblivious adversaries in cellular networks of reuse distance* $k \geq 2$.

The term $\epsilon(k)$ is quite large for small values of k. For any value of reuse distance k, the exact bound implied by Theorem 38.12 is significantly smaller than the corresponding lower bound of deterministic algorithms.

Suitable colorings could be used in wireless networks with different interference graphs. However, for more general interference graphs, irregularities may be difficult to handle. An important class of wireless networks that also contains the cellular networks of reuse distance 2 is the class of planar graphs. These graphs are known to be 4-colorable. Using such a coloring and applying the "classify and randomly select" paradigm, we obtain a 4-competitive online algorithm [20]. No better bounds are known for this case.

However, planarity is a property broad enough to accurately model the wireless networks deployed in practice. In Reference 5, we study wireless networks in which the interference graph has small degree. This case can model the irregularity in the distribution and shape of cells in areas where establishing a regular infrastructure is difficult owing to environmental factors (e.g., mountains, etc.) or just economically infeasible to achieve. In Reference 5, we present the analysis of p-RANDOM in networks with interference graphs of maximum degree 3 or 4. For suitable values of parameter p, the algorithm achieves competitive ratios that cannot be attained either by deterministic algorithms or by "classify and randomly select" algorithms, and, furthermore, this result extends to wireless networks that support many frequencies. The analysis based on the amortized benefit can be further generalized to show bounds on the competitiveness of p-Random on arbitrary interference graphs beating the lower bounds of deterministic algorithms.

Concerning lower bounds on the competitiveness of randomized algorithms, Yao's principle is the main tool to extend the result presented in Section 38.5 to the networks discussed above. By extending the construction of Theorem 38.11 to cellular networks of reuse distance $k \geq 6$, we obtain a lower bound of 2.5 [8]. A different construction in Reference 6 yields a 2.086 lower bound for randomized call control in planar networks. The results of Bartal et al. [4] show limitations on the power of randomized call control algorithms for networks with more general interference graphs.

Definitely, there are still interesting issues to investigate. First, the upper and lower bounds discussed in this chapter still have small gaps. Determining the best possible competitiveness of call control that can be achieved in cellular networks is a challenging task (especially for cellular networks of small reuse distance). Furthermore, our study is mostly of a combinatorial nature. The assumption that calls have equal duration is not always realistic in practice. But, under the classical competitive analysis model, time seems to be a hard opponent to beat. Introducing durations to the calls would imply meaningless (i.e., extremely high) bounds on the competitiveness of call control, even in very simple scenario; this is a rather disappointing result that mainly shows the limitations of competitive analysis. Considering statistical adversaries [3] that construct sequences of calls according to probability distributions known to the algorithm seems to provide a realistic model in order to handle calls with different duration.

Finally, we point out that we have focused on the optimization of communication between users and base stations. Extending the notion of calls to direct connections between pairs of users is worth investigating. By applying similar arguments to those presented in this chapter, we obtain lower bounds of 6 and 3.5 for deterministic and randomized online call control algorithms, respectively. Randomization is definitely helpful in this setting as well, since our preliminary results demonstrate that the competitiveness of p-Random is significantly smaller than 5 for single-frequency wireless cellular networks. Unfortunately, the techniques of Reference 1 for extending the single-frequency case to the case of many frequencies are not applicable to this model. Here, new analysis techniques deserve some attention.

References

[1] B. Awerbuch, Y. Azar, A. Fiat, S. Leonardi, and A. Rosen. On-line competitive algorithms for call admission in optical networks. *Algorithmica*, Vol. 31(1), pp. 29–43, 2001.

[2] B. Awerbuch, Y. Bartal, A. Fiat, and A. Rosen. Competitive non-preemptive call control. In *Proc. of the 5th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA '94)*, pp. 312–320, 1994.

[3] A. Borodin and R. El-Yaniv. *Online Computation and Competitive Analysis.* Cambridge University Press, Cambridge, UK, 1998.

[4] Y. Bartal, A. Fiat, and S. Leonardi. Lower bounds for on-line graph problems with applications to on-line circuit and optical routing. In *Proc. of the 28th Annual ACM Symposium on Theory of Computing (STOC '96)*, pp. 531–540, 1996.

[5] I. Caragiannis, C. Kaklamanis, and E. Papaioannou. Randomized call control in sparse wireless cellular networks. In *Proc. of the 8th International Conference on Advances in Communications and Control (COMCON 01)*, pp. 73–82, 2001.

[6] I. Caragiannis, C. Kaklamanis, and E. Papaioannou. Efficient on-line frequency allocation and call control in cellular networks. *Theory of Computing Systems*, Vol. 35, pp. 521–543, 2002. Preliminary results in *ACM SPAA '00* and *IPDPS '01*.

[7] I. Caragiannis, C. Kaklamanis, and E. Papaioannou. Simple on-line algorithms for call control in cellular networks. In *Proc. of the 1st Workshop on Approximation and On-line Algorithms (WAOA '03)*, LNCS 2909, Springer, pp. 67–80, 2003.

[8] I. Caragiannis, C. Kaklamanis, and E. Papaioannou. New bounds on the competitiveness of randomized online call control in cellular networks. In *Proc. of the Euro-Par '05*, LNCS 3648, Springer, pp. 1089–1099, 2005.

[9] D. Dimitrijević and J. Vučetić. Design and performance analysis of algorithms for channel allocation in cellular networks. *IEEE Transactions on Vehicular Technology*, Vol. 42(4), pp. 526–534, 1993.

[10] W. K. Hale. Frequency assignment: Theory and applications. *Proceedings of the IEEE*, 68(12), pp. 1497–1514, 1980.

[11] J. Janssen, K. Kilakos, and O. Marcotte. Fixed preference frequency allocation for cellular telephone systems. *IEEE Transanctions on Vehicular Technology*, Vol. 48(2), pp. 533–541, 1999.

[12] J. Janssen, D. Krizanc, L. Narayanan, and S. Shende. Distributed on-line frequency assignment in cellular networks. *Journal of Algorithms*, Vol. 36(2), pp. 119–151, 2000.

[13] S. Jordan and E. Schwabe. Worst-case performance of cellular channel allocation policies. *ACM Baltzer Journal on Wireless Networks*, Vol. 2(4), pp. 265–275, 1996.

[14] S. Kim and S. L. Kim. A Two-phase algorithm for frequency assignment in cellular mobile systems. *IEEE Transactions on Vehicular Technology*, Vol. 43(3), pp. 542–548, 1994.

[15] S. Leonardi, A. Marchetti-Spaccamela, A. Prescuitti, and A. Rosen. On-line randomized call-control revisited. In *Proccedings of the 9th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA '98)*, pp. 323–332, 1998.

[16] R. Motwani and B. Raghavan. *Randomized Algorithms*. Cambridge University Press, Cambridge, UK, 1995.

[17] C. McDiarmid and B. Reed. Channel assignment and weighted coloring. *Networks*, Vol. 36, pp. 114–117, 2000.

[18] L. Narayanan and S. Shende. Static frequency assignment in cellular networks. In *Algorithmica*, Vol. 29(3), pp. 396–409, 2001.

[19] L. Narayanan and Y. Tang. Worst-case analysis of a dynamic channel assignment strategy. *Discrete Applied Mathematics*, Vol. 140(1-3), pp. 115–141, 2004.

[20] G. Pantziou, G. Pentaris, and P. Spirakis. Competitive call control in mobile networks. *Theory of Computing Systems*, Vol. 35(6), pp. 625–639, 2002.

[21] P. Raymond. Performance analysis of cellular networks. *IEEE Transactions on Communications*, Vol. 39(12), pp. 1787–1793, 1991.

[22] P.-J. Wan and L. Liu. Maximal throughput in wavelength-routed optical networks. Multichannel optical networks: Theory and practice, *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, AMS, Vol. 46, pp. 15–26, 1998.

[23] A. C. Yao. Probabilistic computations: Towards a unified measure of complexity. In *Proc. of the 17th Annual Symposium on Foundations of Computer Science (FOCS '77)*, pp. 222–227, 1977.