Simple On-Line Algorithms for Call Control in Cellular Networks*

Ioannis Caragiannis, Christos Kaklamanis, and Evi Papaioannou

Computer Technology Institute and
Dept. of Computer Engineering and Informatics
University of Patras, 26500 Rio, Greece

Abstract. We address an important communication issue in wireless cellular networks that utilize Frequency Division Multiplexing (FDM) technology. In such networks, many users within the same geographical region (cell) can communicate simultaneously with other users of the network using distinct frequencies. The spectrum of the available frequencies is limited; thus, efficient solutions to the call control problem are essential. The objective of the call control problem is, given a spectrum of available frequencies and users that wish to communicate, to maximize the number of users that communicate without signal interference. We consider cellular networks of reuse distance $k \geq 2$ and we study the on-line version of the problem using competitive analysis.

In cellular networks of reuse distance 2, the previously best known algorithm that beats the lower bound of 3 on the competitiveness of deterministic algorithms works on networks with one frequency, achieves a competitive ratio against oblivious adversaries which is between 2.469 and 2.651, and uses a number of random bits at least proportional to the size of the network. We significantly improve this result by presenting a series of simple randomized algorithms that have competitive ratios smaller than 3, work on networks with arbitrarily many frequencies, and use only a constant number of random bits or a comparable weak random source. The best competitiveness upper bound we obtain is 7/3.

In cellular networks of reuse distance $k > 2$, we present simple randomized on-line call control algorithms with competitive ratios which significantly beat the lower bounds on the competitiveness of deterministic ones and use only $O(\log k)$ random bits. Furthermore, we show a new lower bound on the competitiveness of on-line call control algorithms in cellular networks of reuse distance $k \geq 5$.

1 Introduction

In this paper we study frequency spectrum management issues in wireless networks. Due to the rapid growth of wireless communications and the limitations and cost of the frequency spectrum, such issues are very important nowadays.

* This work was partially funded by the European Union under IST FET Project ALCOM–FT, IST FET Project CRESCCO and RTN Project ARACNE.

R. Solis-Oba and K. Jansen (Eds.): WAOA 2003, LNCS 2909, pp. 67–80, 2004.

We consider wireless networks in which base stations are used to build the required infrastructure. In such systems, the architectural approach used is the following. A geographical area in which communication takes place is divided into regions. Each region is the calling area of a base station. Base stations are connected via a high speed network.When a user A wishes to communicate with some other user B, a path must be established between the base stations of the regions in which the users A and B are located. Then communication is performed in three steps: (a) wireless communication between A and its base station, (b) communication between the base stations, and (c) wireless communication between B and its base station. Thus, the transmission of a message from A to B first takes place between A and its base station, the base station of A sends the message to the base station of B which will transmit it to B. 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 wireless 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 [6, 7] is the one shown in the left part of Figure 1. All regions are regular hexagons (cells) of the same size. This shape results from the 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. Due to the shape of the regions, we call these networks cellular wireless networks.

Many users of the same region can communicate simultaneously with their base station. This can be achieved via 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. Since the spectrum of available frequencies is limited, important engineering problems related to the efficient reuse of frequencies arise. Signal interference usually manifests itself when the same frequency is assigned to users located in the same or adjacent cells. Alternatively, in this case, we may say that the cellular network has *reuse distance* 2. By generalizing this parameter, we obtain cellular networks of reuse distance k in which signal interference between users that have been assigned the same frequency is avoided only if the users are located in cells with distance at least k.

Signal interference in cellular networks can be represented by an *interference graph* G. Vertices of the graph correspond to cells and an edge (u, v) in the graph indicates that the assignment of the same frequency to two users lying at the cells corresponding to nodes u and v will cause signal interference. The interference graph of a cellular network of reuse distance 2 is depicted in the right part of Figure 1. If the assumption of uniform distribution of identical base stations does not hold, arbitrary interference graphs can be used to model the underlying network.

In this paper we study the *call control* problem which is defined as follows: Given users that wish to communicate with their base station, the *call control* problem on a network that supports a spectrum of w available frequencies is to

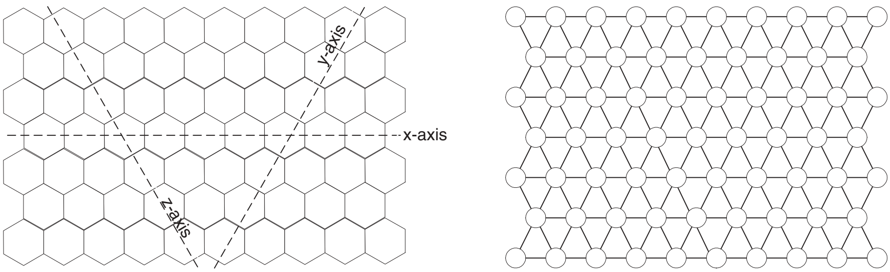

Fig. 1. A cellular network and the interference graph if the reuse distance is 2.

assign frequencies to users so that at most w frequencies are used in total, signal interference is avoided, and the number of users served is maximized.

We assume that calls corresponding to users that wish to communicate appear in the cells of the network in an on–line manner. When a call arrives, a call control algorithm decides either to accept the call (assigning a frequency to it), or to reject it. Once a call is accepted, it cannot be rejected (preempted). Furthermore, the frequency assigned to the call cannot be changed in the future. We assume that all calls have infinite duration; this assumption is equivalent to considering calls of the same duration.

Competitive analysis [11] has been used for evaluating the performance of on–line algorithms for various problems. In our setting, given a sequence of calls, the performance of an on–line algorithm A is compared to the performance of the optimal algorithm OPT.

Let $B(\sigma)$ be the benefit of the on–line algorithm A on the sequence of calls σ, i.e. the set of calls of σ accepted by A and $O(\sigma)$ the benefit of the optimal algorithm. If A is a deterministic algorithm, we define its competitive ratio (or competitiveness) as $\max_\sigma \frac{|O(\sigma)|}{|B(\sigma)|}$, where the maximum is taken over all possible sequences of calls. If A is a randomized algorithm, we define its competitive ratio as $\max_\sigma \frac{|O(\sigma)|}{\mathcal{E}[|B(\sigma)|]}$, where $\mathcal{E}[|B(\sigma)|]$ is the expectation of the number of calls accepted by A, and the maximum is taken over all possible sequences of calls.

Usually, we compare the performance of deterministic algorithms against *off–line adversaries*, i.e. adversaries that have knowledge of the behaviour of the deterministic algorithm in advance. In the case of randomized algorithms, we consider *oblivious adversaries* whose knowledge is limited to the probability distribution of the random choices of the randomized algorithm.

The static version of the call control problem is very similar to the famous maximum independent set problem. The on–line version of the problem is studied in [1–4, 8, 10]. [1], [2], and [8] study the call control problem in the context of optical networks. Pantziou et al. [10] present upper bounds for networks with planar and arbitrary interference graphs. Usually, competitive analysis of call control focuses on networks supporting one frequency. Awerbuch et al. [1] present a simple way to transform algorithms designed for one frequency to algorithms for arbitrarily many frequencies with a small sacrifice in competitiveness (see also [5, 12]). Lower bounds for call control in arbitrary networks are presented in [3].

The greedy algorithm is probably the simplest on-line algorithm. When a call arrives, the greedy algorithm seeks for the first available frequency. If such a frequency exists, the algorithm accepts the call assigning this frequency to it, otherwise, the call is rejected. In general, Pantziou et al. [10] show that this algorithm has competitive ratio equal to the degree of the interference graph and no better in general. The greedy algorithm is optimal within the class of deterministic on-line call control algorithms.

Simple randomized algorithms can be defined using the "classify and randomly select" paradigm [1, 2, 10]. Such algorithms use a coloring of the underlying interference graph, randomly select a color out of the colors used, and execute the greedy algorithm in the cells colored with the selected color, ignoring (i.e., rejecting) calls in all other cells. The competitive ratio achieved in this way, against oblivious adversaries, is equal to the number of colors used in the coloring of the interference graph.

In cellular networks of reuse distance 2, the greedy algorithm is 3-competitive against off-line adversaries, in the case of one frequency. Slightly worse competitiveness bounds can be proved in the case of arbitrarily many frequencies using the techniques of [1, 5, 12]. In [4], using similar arguments with those of [10], it was observed that no deterministic on-line call control algorithm in cellular networks of reuse distance 2 can be better than 3-competitive against off-line adversaries. Applying the "classify and randomly select" paradigm using a 3-coloring of the interference graph, we obtain a 3-competitive randomized algorithm even in the case of arbitrarily many frequencies. Observe that this algorithm uses a very weak random source which equiprobably selects one out of three distinct objects.

The authors in [4] describe algorithm p–Random, an intuitive on–line randomized call control algorithm for networks that support one frequency. They present upper and lower bounds on the competitive ratio of the algorithm as functions of parameter p and, by optimizing these functions, they prove that, for some value of p, the competitive ratio of algorithm p-Random against oblivious adversaries is between 2.469 an 2.651. The analysis of algorithm p-Random in [4] applies only to cellular networks with one frequency but it indicates that randomization helps to beat the deterministic upper bounds. However, the number of random bits used by the algorithm may be proportional to the size of the network. The best known lower bound on the competitive ratio of any randomized call control algorithm in cellular networks of reuse distance 2 is 1.857 [4].

Results on on-line call control in cellular networks of reuse distance $k > 2$ are only implicit in previous work. The greedy algorithm has competitive ratio 4 and 5 in cellular networks of reuse distance $k \in \{3, 4, 5\}$ and $k \geq 6$, respectively, which support one frequency. This is due to the fact that the acceptance of a non-optimal call may cause the rejection of at most 4 and 5 optimal calls, respectively. These competitive ratios are the best possible that can be achieved by deterministic algorithms. Using the techniques of [1, 5, 12], it can be shown that, in the case of arbitrarily many frequencies, the greedy algorithm has competitive ratio at most 4.521 and at most 5.517 in cellular networks of reuse distance $k \in \{3, 4, 5\}$ and $k \geq 6$, respectively.

Furthermore, applying the "classify and randomly select" paradigm using an efficient coloring of the interference graph of cellular networks of distance reuse $k > 2$ would give randomized on-line algorithms with competitive ratio $\Omega(k^2)$. Even in the case of $k = 3$, the competitive ratio we obtain in this way is 7.

In this paper, we improve previous results on the competitiveness of on-line call control algorithms in cellular networks. We present algorithms based on the "classify and randomly select" paradigm which use new colorings of the interference graph. These algorithms use a small number of random bits, and have small competitive ratios against oblivious adversaries even in the case of arbitrarily many frequencies. In particular, in cellular networks of reuse distance 2, we significantly improve the best known competitiveness bounds achieved by algorithm p-RANDOM by presenting a series of simple randomized algorithms that have smaller competitive ratios, work on networks with arbitrarily many frequencies, and use only a constant number of random bits or a comparable weak random source. The best competitiveness upper bound we obtain is 7/3. In cellular networks of reuse distance $k > 2$, we present simple randomized on-line call control algorithms with competitive ratios which significantly beat the lower bounds on the competitiveness of deterministic algorithms and use only $O(\log k)$ random bits. For any $k > 2$, the competitive ratio we achieve is strictly smaller than 4. Furthermore, we show a new lower bound of 25/12 on the competitiveness, against oblivious adversaries, of on-line call control algorithms in cellular networks of reuse distance $k \geq 5$.

The rest of the paper is structured as follows. We present a simple 8/3-competitive randomized algorithm in Section 2. Simpler algorithms with better competitive ratios (including the algorithm with competitive ratio 7/3) are presented in Section 3. The upper bounds for cellular networks of reuse distance $k > 2$ are presented in Section 4. A Our new lower bound for cellular networks of reuse distance $k \geq 5$ can be found in Section 5.

2 A Simple 8/3-Competitive Algorithm

In this section we present algorithm CRS-A, a simple randomized on-line algorithm for call control in cellular networks of reuse distance 2. The algorithm works in networks with one frequency and achieves a competitive ratio against oblivious adversaries which is similar (but slightly inferior) to that which has been 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-axis with either the colors 0 and 1 or the colors 2 and 3, coloring the cells in the same y-axis with either the colors 0 and 2 or the colors 1 and 3, and coloring the cells in the same z-axis with either the colors 0 and 3 or the colors 1 and 2. Such a coloring is depicted in the left part of Figure 2.

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.

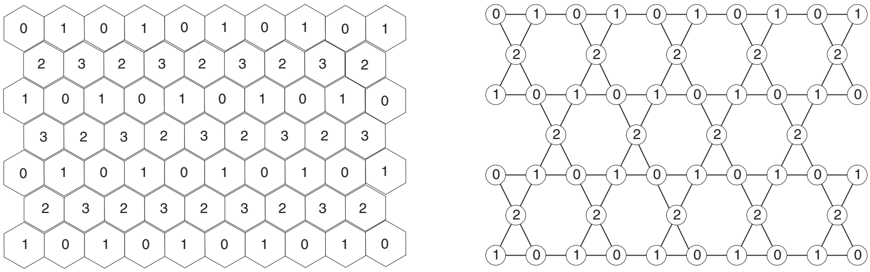

Fig. 2. The 4-coloring used by algorithm CRS-A and the corresponding subgraph of the interference graph induced by the nodes not colored with color 3.

Theorem 1. *Algorithm CRS-A in cellular networks of reuse distance* 2 *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 which 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 which 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 $\mathcal{E}[|O'|] \geq \frac{3}{4}|O|$. Now let B be the set of calls accepted by algorithm CRS-A, i.e., 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 on-line algorithms with improved competitive ratios that use slightly stronger random sources and work on networks with arbitrarily many frequencies.

3 Improved Algorithms

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., [1], [10]).

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

Lemma 1. *Consider a network with interference graph $G = (V, E)$ which 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 $s_0, s_1, ..., s_{\nu-1} \subseteq X$ and an integer λ such that*

- *each color of X belongs to at least λ different sets of the sets $s_0, s_1, ..., s_{\nu-1}$, and*
- *for $i = 0, 1, ..., \nu-1$, 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 on-line 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 which supports w frequencies and the randomized on-line algorithm working as follows. The algorithm randomly selects one out of the ν color classes $s_0, ..., s_{\nu-1}$ 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 two randomized on-line algorithms for call control in cellular networks of reuse distance 2, namely CRS-B and CRS-C, 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-axis 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 3. Also, define $s_i = \{i, (i+1) \bmod 5\}$, for $i = 0, 1, ..., 4$. Observe that, for each $i = 0, 1, ..., 4$, 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 $(i+4) \bmod 5$, i.e., colors not belonging to s_i. Thus, the coloring χ together with the color classes s_i satisfy the conditions of Lemma 1 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 1. We obtain the following.

Fig. 3. The 5-coloring used by algorithm CRS-B and the 7-coloring used by algorithm CRS-C. The skewed cells are those colored with the colors in set s_0.

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

Now consider a coloring of the cells with seven colors $0, 1, ..., 6$ such that for each for each cell colored with color i (for $i = 0, ..., 6$), its two adjacent cells in the same x-axis are colored with the colors $(i-1) \bmod 7$ and $(i+1) \bmod 7$, while its two adjacent cells in the same z-axis are colored with colors $(i-3) \bmod 7$ and $(i+3) \bmod 7$. Such a coloring is depicted in the right part of Figure 3. Also, define $s_i = \{i, (i+1) \bmod 7, (i+3) \bmod 7\}$, for $i = 0, 1, ..., 6$. Observe that, for each $i = 0, 1, ..., 6$, 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 $(i+6) \bmod 7$, i.e., colors not belonging to s_i. Thus, the coloring χ together with the color classes s_i satisfy the conditions of Lemma 1 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 1. We obtain the following.

Theorem 3. *Algorithm CRS-C in cellular networks with reuse distance* 2 *is* 7/3-*competitive against oblivious adversaries.*

The two algorithms above (CRS-B and CRS-C) make use of a random source which equiprobably selects one out of an odd number of distinct objects. If we only have a number of fair coins (random bits) available, we can design algorithms with small competitive ratios by combining the algorithms above. For example, using 6 random bits, we may construct the following algorithm. We use integers $0, 1, ..., 63$ to identify each of the 63 outcomes of the 6 random bits. For an outcome $i \in \{0, ..., 49\}$, the algorithm executes algorithm CRS-B using color class $s_{i \bmod 5}$, and for an outcome $i \in \{50, ..., 63\}$, the algorithm executes algorithm CRS-C using color class $s_{(i-50) \bmod 7}$. It can be easily seen that this algorithm has competitive ratio $32/13 \approx 2.462$ against oblivious adversaries, since its expected benefit is at least $50/64 \cdot 2/5 + 14/64 \cdot 3/7 = 13/32$ times the optimal benefit. Similarly, using 8 random bits, we obtain an on-line algorithm with competitive ratio $64/27 \approx 2.37$. We can generalize this idea, and, for sufficiently small $\epsilon > 0$, we can construct an algorithm which uses $t = O(\log 1/\epsilon)$ random bits, and, on $2^t \bmod 7$ of the 2^t outcomes, it does nothing, while the rest of the outcomes are assigned to executions of algorithm CRS-C. In this way, we obtain the following.

Corollary 1. *For any $\epsilon > 0$, there exists an on-line randomized call-control algorithm for cellular networks with reuse distance* 2 *that uses $O(\log 1/\epsilon)$ random bits and has competitive ratio at most $7/3 + \epsilon$ against oblivious adversaries.*

4 Cellular Networks with Reuse Distance $k > 2$

For cellular networks with reuse distance $k > 2$, we present algorithm CRS-k which is based on the "classify and randomly select" paradigm. Algorithm CRS-k uses the following coloring of the interference graph of a cellular network with reuse distance k. Cells are colored with the colors $0, 1, ..., 3k^2 - 3k$ such that for any cell colored with color i, its adjacent cells in the x-axis are colored with colors $(i-1) \bmod (3k^2-3k+1)$ and $(i+1) \bmod (3k^2-3k+1)$, while its adjacent cells in the z-axis are colored with colors $(i-3(k-1)^2) \bmod (3k^2-3k+1)$ and $(i+3(k-1)^2) \bmod (3k^2-3k+1)$.

For odd k, for $i = 0, 1, ..., 3k^2 - 3k$, the color class s_i contains the following colors. For $j = 0, 1, ..., \frac{k-1}{2}$, it contains the colors $(i + 3j(k-1)^2 - j) \bmod (3k^2 - 3k+1)$, ..., $(i + 3j(k-1)^2 + \frac{k-1}{2}) \bmod (3k^2 - 3k + 1)$, and for $j = \frac{k+1}{2}, ..., k-1$, it contains the colors $(i + \frac{3(k-1)^3}{2} + 3(j - \frac{k-1}{2})(k-1)^2 - \frac{k-1}{2}) \bmod (3k^2 - 3k+1)$, ..., $(i + \frac{3(k-1)^3}{2} + 3(j - \frac{k-1}{2})(k-1)^2 + k - 1 - j) \bmod (3k^2 - 3k + 1)$.

For even k, for $i = 0, 1, ..., 3k^2 - 3k$, the color class s_i contains the following colors. For $j = 0, 1, ..., \frac{k}{2} - 1$, it contains the colors $(i + 3j(k-1)^2 - j) \bmod (3k^2 - 3k+1), ..., (i+3j(k-1)^2 + \frac{k}{2}) \bmod (3k^2-3k+1)$, and for $j = \frac{k}{2}, ..., k-1$, it contains the colors $(i + 3(\frac{k}{2} - 1)(k-1)^2 + 3(j - \frac{k}{2} + 1)(k-1)^2 - \frac{k}{2} + 1) \bmod (3k^2 - 3k + 1)$, ..., $(i + 3(\frac{k}{2} - 1)(k-1)^2 + 3(j - \frac{k}{2} + 1)(k-1)^2 + k - 1 - j) \bmod (3k^2 - 3k + 1)$.

Note that, for $k = 2$, we obtain the coloring used by algorithm CRS-C. Examples of the coloring for $k = 3$ and $k = 4$ as well as the cells colored with colors from the color class s_0 are depicted in Figure 4.

We can show the following two lemmas. Formal proofs will be given in the final version of the paper.

Lemma 2. *Let $k > 2$ and G be the interference graph of a cellular network of distance reuse k. Consider the coloring of G used by algorithm CRS-k and the color classes s_i, for $i = 0, 1, ..., 3k^2 - 3k$. For any color j such that $0 \leq j \leq 3k^2 - 3k$, the number of different color classes s_i that color j belongs to is $\frac{3k^2}{4}$ if k is even, and $\frac{3k^2+1}{4}$ if k is odd.*

Lemma 3. *Let $k > 2$ and G be the interference graph of a cellular network of distance reuse k. Consider the coloring of G used by algorithm CRS-k and the color classes s_i, for $i = 0, 1, ..., 3k^2 - 3k$. For $i = 0, 1, ..., 3k^2 - 3k$, each connected component of the subgraph of G induced by the nodes of G colored with colors in s_i is a clique.*

Thus, the colorings and the color classes described satisfy the condition of Lemma 1 with $\lambda = 3k^2 - 3k + 1$ and $\nu = \frac{3k^2}{4}$ if k is even, and $\nu = \frac{3k^2+1}{4}$ if k is odd. By Lemma 1, we obtain the following.

Theorem 4. *The competitive ratio against oblivious adversaries of algorithm CRS-k in cellular networks with reuse distance $k \geq 2$ is $4\left(1 - \frac{3k-1}{3k^2}\right)$ if k is even, and $4\left(1 - \frac{3k}{3k^2+1}\right)$ if k is odd.*

Algorithm CRS-k in cellular networks of reuse distance k uses a random source which equiprobably selects one among $3k^2 - 3k + 1$ distinct objects. By applying ideas we used in the previous section, we can achieve similar competitiveness bounds by algorithms that use random bits.

Corollary 2. *For any $\epsilon > 0$, there exists an on-line randomized call-control algorithm for cellular networks of reuse distance k, that uses $O(\log 1/\epsilon + \log k)$ random bits and has competitive ratio at most $4\left(1 - \frac{3k-1}{3k^2}\right) + \epsilon$ if k is even, and $4\left(1 - \frac{3k}{3k^2+1}\right) + \epsilon$ if k is odd, against oblivious adversaries.*

5 Lower Bounds

In previous work ([4], implicitly in [10]), it has been observed that no deterministic on-line call control algorithm can have a competitive ratio better than 3 against off-line adversaries. We can easily extend this lower bound and obtain

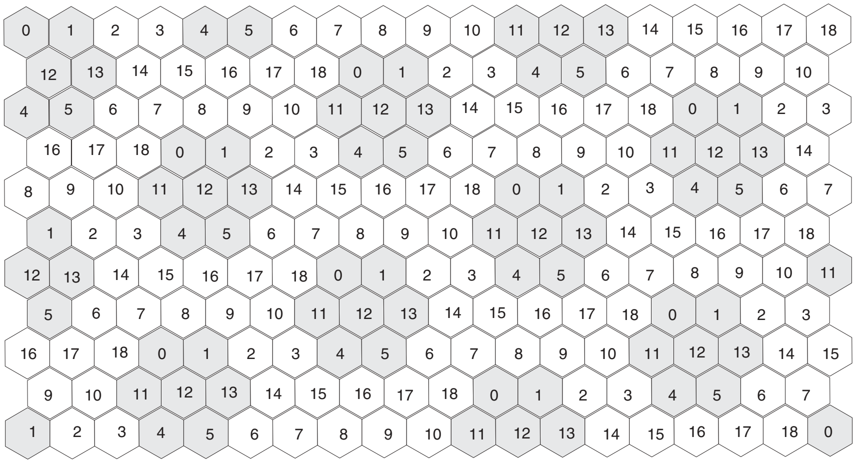

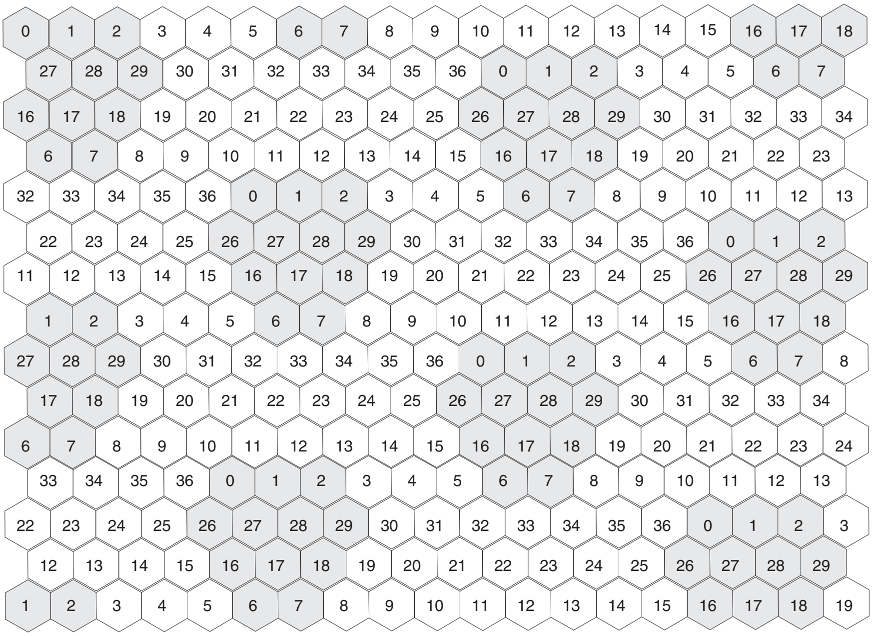

Fig. 4. Examples of the colorings used by the algorithms CRS-3 and CRS-4. The skewed cells are those colored with the colors in set s_0.

lower bounds of 4 and 5 on the competitiveness of deterministic on-line call control algorithms in cellular networks of reuse distance $k \in \{3, 4, 5\}$ and $k \geq 6$, respectively.

Hence, 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 [13] (see also [9]), we prove a lower bound on the competitive ratio, against oblivious adversaries, of any randomized algorithm in cellular networks with reuse distance $k \geq 5$. Again, we consider networks that support one frequency; our lower bound can be trivially extended to networks that support multiple frequencies. In our proof, we use the following lemma.

Lemma 4 (Minimax Principle [9]). *Given a probability distribution $\mathcal{P}$ over sequences of calls σ, denote by $\mathcal{E}_\mathcal{P}[B_A(\sigma)]$ and $\mathcal{E}_\mathcal{P}[B_{OPT}(\sigma)]$ the expected benefit of a deterministic algorithm A and the optimal off–line 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*

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

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, i.e. $c_A^\mathcal{P} \leq c_{A_R}$.

Our lower bound is the following.

Theorem 5. *No randomized call–control algorithm in cellular networks with distance reuse $k \geq 5$ can be better than 25/12–competitive against an oblivious adversary.*

Proof. Consider a cellular network with reuse distance 5 and ten cells v_0, $v_{A,1}$, $v_{A,2}$, $v_{B,1}$, $v_{B,2}$, $v_{B,3}$, $v_{C,1}$, $v_{C,2}$, $v_{C,3}$, and $v_{C,4}$ as shown in Figure 5. We will prove that there exists an adversary $\mathcal{ADV}$ that produces calls in these cells according to a probability distribution $\mathcal{P}$ in such way that no deterministic algorithm can be better than 25/12–competitive under $\mathcal{P}$ even if it knows the probability distribution $\mathcal{P}$ in advance.

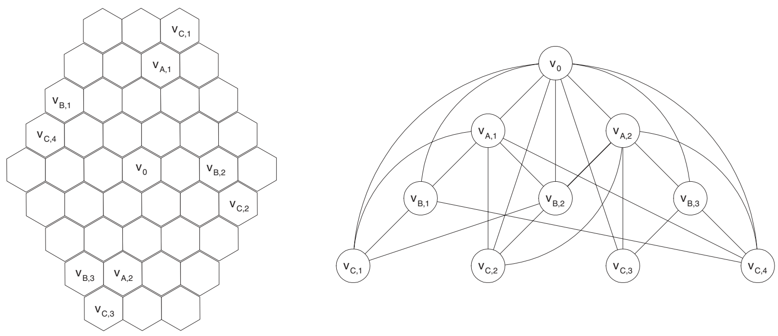

Fig. 5. The cellular network of reuse distance 5 used in the proof of Theorem 9 and the subgraph of the interference graph induced by the nodes corresponding to the ten cells v_0, $v_{A,1}$, $v_{A,2}$, $v_{B,1}$, $v_{B,2}$, $v_{B,3}$, $v_{C,1}$, $v_{C,2}$, $v_{C,3}$, and $v_{C,4}$.

We define the probability distribution $\mathcal{P}$ as follows. First, the adversary produces a call in the cell v_0. Then, it

- either stops, with probability 1/2,
- or does the following, with probability 1/2. It presents two calls, one in the cell $v_{A,1}$ and one in the cell $v_{A,2}$, and
  - either stops, with probability 1/3,
  - or does the following with probability 2/3. It presents three calls, one in the cell $v_{B,1}$, one in the cell $v_{B,2}$, and one in the cell $v_{B,3}$, and
    - either stops, with probability 1/4,
    - or does the following, with probability 3/4. It presents four calls, one in the cell $v_{C,1}$, one in the cell $v_{C,2}$, one in the cell $v_{C,3}$, and one in the cell $v_{C,4}$, and then stops.

Clearly, the benefit of the optimal off-line algorithm is the number of calls presented by the adversary in the last step before stopping the sequence. Thus, the expected benefit of the optimal off–line algorithm on sequences of calls generated according to $\mathcal{P}$ is

$$\mathcal{E}_{\mathcal{P}}[B_{OPT}(\sigma)] = 1 \cdot \frac{1}{2} + 2 \cdot \frac{1}{2} \cdot \frac{1}{3} + 3 \cdot \frac{1}{2} \cdot \frac{2}{3} \cdot \frac{1}{4} + 4 \cdot \frac{1}{2} \cdot \frac{2}{3} \cdot \frac{3}{4} = \frac{25}{12}.$$

Let A be a deterministic call control algorithm that runs on the calls produced by $\mathcal{ADV}$. First, we observe that no algorithm could gain by accepting one of the two calls presented in cells $v_{A,1}$ and $v_{A,2}$ or by accepting one or two of the three calls presented in cells $v_{B,1}$, $v_{B,2}$ and $v_{B,3}$. Hence, we may assume that the algorithm either accepts all calls presented at a step or rejects them all.

Consider t executions of the algorithm on t sequences produced according to the probability distribution $\mathcal{P}$. Let q_0 be the number of executions in which A accepts the call produced in cell v_0, q_1 the number of executions in which A accepts both calls in cells $v_{A,1}$ and $v_{A,2}$, and q_2 the number of executions in which the algorithm accepts the three calls in cells $v_{C,1}$, $v_{C,2}$, and $v_{C,3}$.

The expected number of executions in which the algorithm does not accept the call in cell v_0 and the adversary produces calls in cells $v_{A,1}$ and $v_{A,2}$ is $\frac{1}{2}(t-q_0)$. Hence, the expected number of executions in which the algorithm does not accept the calls in cells v_0, $v_{A,1}$, and $v_{A,2}$ and the adversary produces calls in cells $v_{B,1}$, $v_{B,2}$, and $v_{B,3}$ is $\frac{2}{3}\left(\frac{1}{2}(t-q_0)-q_1\right)$ and the expected number of executions in which the algorithm does not accept the calls in cells v_0, $v_{A,1}$, $v_{A,2}$, $v_{B,1}$, $v_{B,2}$, and $v_{B,3}$ and the adversary produces calls in cells $v_{C,1}$, $v_{C,2}$, $v_{C,3}$, and $v_{C,4}$ is $\frac{3}{4}\left(\frac{2}{3}\left(\frac{1}{2}(t-q_0)-q_1\right)-q_2\right)$. Thus,

$$\mathcal{E}_{\mathcal{P}}[B_A(\sigma)] \leq \frac{q_0 + 2q_1 + 3q_2 + 4\frac{3}{4}\left(\frac{2}{3}\left(\frac{1}{2}(t-q_0)-q_1\right)-q_2\right)}{t} = 1$$

and $c_A^{\mathcal{P}} \geq 25/12$. By Lemma 4, we obtain that this is a lower bound on the competitiveness of any randomized algorithm against oblivious adversaries. □

References

1. B. Awerbuch, Y. Azar, A. Fiat, S. Leonardi, and A. Rosen. Competitive 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, A. Rosen. Competitive Non–Preemptive Call Control. In *Proccedings of the 5th Annual ACM–SIAM Symposium on Discrete Algorithms (SODA '94)*, pp. 312–320, 1994.
3. 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)*, 1996.
4. 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.
5. T. Erlebach and K. Jansen. The Maximum Edge-Disjoint Paths Problem in Bidirected Trees. *SIAM Journal on Discrete Mathematics*, Vol. 14(3), pp. 326-355, 2001.
6. W.K. Hale. Frequency Assignment: Theory and Applications. In *Proceedings of the IEEE*, 68(12), pp. 1497–1514, 1980.
7. 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.
8. 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.
9. R. Motwani and B. Raghavan. Randomized Algorithms. *Cambridge University Press*, 1995.
10. G. Pantziou, G. Pentaris, and P. Spirakis. Competitive Call Control in Mobile Networks. *Theory of Computing Systems*, Vol. 35(6), pp. 625-639, 2002.
11. D. Sleator and R.E. Tarjan. Amortized Efficiency of List Update and Paging Rules. *Communications of Association of Computing Machinery* 28, pp. 202–208, 1985.
12. 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.
13. A. C. Yao. Probabilistic Computations: Towards a Unified Measure of Complexity. In *Proceedings of the 17th Annual Symposium on Foundations of Computer Science (FOCS '77)*, pp. 222–227, 1977.