

New Bounds on the Competitiveness of Randomized Online Call Control in Cellular Networks*

Ioannis Caragiannis, Christos Kaklamanis, and Evi Papaioannou

Research Academic Computer Technology Institute and
Dept. of Computer Engineering and Informatics
University of Patras, 26500 Rio, Greece
{caragian,kakl,papaioan}@ceid.upatras.gr

Abstract. We address the call control problem 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 available frequency spectrum is limited; hence, its management should be done efficiently. The objective of the call control problem is, given a spectrum of available frequencies and users that wish to communicate in a cellular network, to maximize the number of users that communicate without signal interference. We study the online version of the problem in cellular networks using competitive analysis and present new upper and lower bounds.

1 Introduction

In this paper we study frequency spectrum management issues in wireless networks. 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 where 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. 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.

Network Model. The network topology usually adopted [8, 9] is the one shown in the left part of Figure 1. All regions are regular hexagons (cells) of the same

* This work was partially funded by the European Union under the IST FET Project CRESCCO.

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 of the network 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 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 whose vertices correspond to cells and an edge (u, v) 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.




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

Problem Definition. In this paper we study the *call control* (or call admission) problem which is defined as follows: Given users that wish to communicate, the *call control* problem on a network that supports a spectrum of w available frequencies is to 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 online 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 [4] has been used for evaluating the performance of online algorithms for various problems. In our setting, given a sequence of calls, the performance of an online algorithm A is compared to the performance of the optimal algorithm OPT . Let $B(\sigma)$ be the benefit of the online 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. We define the competitive ratio or competitiveness of an algorithm A 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. This definition applies to both deterministic and randomized algorithms. Usually, we compare the performance of deterministic algorithms against *off-line adversaries*, i.e. adversaries that have knowledge of the behavior 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.

Related Work. The static version of the call control problem generalizes the famous maximum independent set problem. The online version of the problem is studied in [1–3, 5, 7, 10, 12]. [1], [2], [7] and [10] study the call control problem in the context of optical networks. Pantziou et al. [12] 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 [7] and [13]). Lower bounds for call control in arbitrary networks are presented in [3].

The greedy algorithm is probably the simplest online algorithm. It considers frequencies as positive integers. When a call arrives, it seeks for the smallest available frequency. If such a frequency exists, the algorithm accepts the call assigning this frequency to it, otherwise, the call is rejected. As observed in [5], this algorithm has competitive ratio equal to the size of the maximum independent set in the neighborhood of any node of the interference graph (see also [12]). The competitive ratio of the greedy algorithm is a lower bound on the competitiveness of every deterministic algorithm. In particular, this gives lower bounds of 3, 4 and 5 on the competitiveness of every deterministic online call control algorithm in cellular networks of reuse distance $k = 2$, $k \in \{3, 4, 5\}$ and $k \geq 6$, respectively.

The first randomized algorithm with competitive ratio smaller than 3 in cellular networks with reuse distance 2 was presented in [5]. The main drawbacks of this algorithm are that it uses a number of random bits which is proportional to the size of the sequences of calls and, that it works in networks that support only one frequency. By extending the “classify and randomly select” paradigm [1, 2, 12], the authors in [6] present a series of simpler randomized algorithms that use a small number of random bits or comparably weak random sources, and have small competitive ratios even in the case of arbitrarily many frequencies, in

cellular networks of any reuse distance k . The best competitive ratio obtained is $4 - \Omega(\frac{1}{k})$, while the randomness used is the ability to select equiprobably one out of an odd number of distinct objects. The best competitive ratio obtained for $k = 2$ is $7/3$. The best known lower bounds on the competitiveness of randomized algorithms are $13/7$ and $25/12$ for cellular networks of reuse distance $k \geq 2$ and $k \geq 5$, respectively ([5, 6]).

Our Results. In this paper, we present (Section 2) a new online call control algorithm with competitive ratio $16/7$ for cellular networks with reuse distance 2, improving the previous best known upper bound of $7/3$. Our algorithm is based on the “classify and randomly select” paradigm, uses only 4 random bits, and works in networks with arbitrarily many frequencies. Furthermore, we show new lower bounds of 2 and 2.5 on the competitiveness against oblivious adversaries of online call control algorithms in cellular networks of reuse distance $k \geq 2$ and $k \geq 6$, respectively (Section 3). Our new lower bounds improve previous ones for almost all cases of the reuse distance ($k \neq 5$).

2 The Upper Bound

In this section, we present the online algorithm CRS-D achieving a competitive ratio of $16/7$ against oblivious adversaries. The algorithm is based on the “classify and randomly select” paradigm. Such algorithms use a coloring of the cells (i.e., a coloring of the nodes of the interference graph) and a classification of the colors into not necessarily disjoint color classes. The algorithm randomly selects one out of the available color classes and executes the greedy algorithm for calls appearing in cells colored with colors from the selected color class, while it completely ignores (i.e., rejects) the calls appearing in any other cell. The following lemma gives a connection between the coloring of the interference graph and the definition of the color classes and the competitiveness of the “classify and randomly select” algorithm that uses them.

Lemma 1 ([6]). *Consider a network with interference graph $G = (V, E)$ 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 (color classes) $s_0, s_1, \dots, s_{\nu-1} \subseteq X$ and an integer $\lambda \leq \nu$ such that*

- *each color of X belongs to at least λ different color classes, and*
- *for $i = 0, 1, \dots, \nu - 1$, each connected component of the subgraph of G induced by the nodes colored with colors in s_i is a clique,*

then the online call control algorithm which uses the coloring χ and the ν color classes according to the “classify and randomly select” paradigm has competitive ratio ν/λ against oblivious adversaries.

A proof was presented in [6]. The intuition behind the proof is that (1) the algorithm runs the greedy algorithm on a fraction of λ/ν of the cells, (2) the optimal solution of the subsequence defined by these calls has size at least λ/ν times the size of the original optimal solution, and (3) the greedy algorithm computes an optimal solution when applied to the subsequence.

Algorithm CRS-D uses a coloring of the cells with sixteen colors $0, \dots, 15$ 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 $0 \leq i, j \leq 3$ as follows:

$$s_{4i+j} = \{4i + j, 4i + (j + 1) \bmod 4, 4((i + 1) \bmod 4) + j, 4((i + 1) \bmod 4) + (j + 2) \bmod 4, 4((i + 2) \bmod 4) + (j + 1) \bmod 4, 4((i + 2) \bmod 4) + (j + 2) \bmod 4, 4((i + 3) \bmod 4) + (j + 3) \bmod 4\}$$

An example of this coloring is depicted in Figure 2.




Fig. 2. The 16-coloring used by algorithm CRS-D. The grey cells are those colored with colors in the class s_0 .

We now show that the coloring and the color classes used by algorithm CRS-D satisfy the conditions of Lemma 1. Each color $k = 0, 1, \dots, 15$ belongs to 7 of the 16 color classes s_0, s_1, \dots, s_{15} . For any i, j such that $0 \leq i, j \leq 3$, color $4i + j$ belongs to the color classes $4i + j, 4i + (j - 1) \bmod 4, 4((i - 1) \bmod 4) + j, 4((i - 1) \bmod 4) + (j - 2) \bmod 4, 4((i - 2) \bmod 4) + (j - 1) \bmod 4, 4((i - 2) \bmod 4) + (j - 2) \bmod 4$, and $4((i - 3) \bmod 4) + (j - 3) \bmod 4$. 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 $4((i + 2) \bmod 4) + (j + 3) \bmod 4$ which 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 $4((i+2) \bmod 4) + (j+2) \bmod 4$, 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 $4((i+2) \bmod 4) + (j+3) \bmod 4$ which 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 $4((i+3) \bmod 4) + (j+2) \bmod 4$ which 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 1 for $\lambda = 7$ and $\nu = 16$. This yields the following.

Theorem 1. *Algorithm CRS-D for call control in cellular networks with reuse distance 2 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.

3 Lower Bounds

In this section, using the Minimax Principle [14] (see also [11]), we prove new lower bounds on the competitive ratio, against oblivious adversaries, of any randomized algorithm in cellular networks with reuse distance $k \geq 2$. 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 2 (Minimax Principle [11]). *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. $c_A^{\mathcal{P}} \leq c_{A_R}$.

So, in order to prove a lower bound for any randomized algorithm, it suffices to define an adversary which 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 (that may know the probability distribution in advance) is above some value; by Lemma 2, this value will also be a lower bound for any randomized algorithm against oblivious adversaries.

Theorem 2.

(a) No randomized online call-control algorithm can be better than 2-competitive against oblivious adversaries in cellular networks with reuse distance $k \geq 2$.

(b) No randomized online call-control algorithm can be better than 2.5-competitive against oblivious adversaries in cellular networks with reuse distance $k \geq 6$.

Proof. Due to lack of space, we prove only the first part of the theorem here. The proof of the second part which uses similar ideas in a slightly more complicated way will appear in the final version of the paper.

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

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

- On HEADS, it produces two calls in cells v_{00} and v_{01} which are at distance at least k from each other, at most $k - 1$ from v_0 and at least k from v_1 . Then, it stops.
- On TAILS, it produces two calls in cells v_{10} and v_{11} which are at distance at least k from each other, at most $k - 1$ from v_1 and at least k from v_0 . Then, it stops.

Now, consider the set of all possible deterministic algorithms \mathcal{A}_1 working on the sequences produced by $\mathcal{ADV}\text{-}1$. 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 $\mathcal{ADV}\text{-}1$ are at distance at least k from 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 at distance at most $k - 1$ from 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 $\mathcal{ADV}\text{-}2$ works as follows: It first produces two calls in cells v_0 and v_1 which have distance at least k . Then it tosses a fair coin.

- On HEADS, it produces two calls in cells v_{00} and v_{01} which are at distance at least k from each other, at most $k - 1$ from v_0 and at least k from v_1 . Then, it tosses a fair coin.
 - On HEADS, it produces two calls in cells v_{000} and v_{001} which are at distance at least k from each other, at most $k - 1$ from v_0 and v_{00} , and at least k from v_1 and v_{01} . Then, it stops.
 - On TAILS, it produces two calls in cells v_{010} and v_{011} which are at distance at least k from each other, at most $k - 1$ from v_0 and v_{01} and at least k from v_1 and v_{00} . Then, it stops.
- On TAILS, it produces two calls in cells v_{10} and v_{11} which are at distance at least k from each other, at most $k - 1$ from v_1 and at least k from v_0 . Then, it tosses a fair coin.
 - On HEADS, it produces two calls in cells v_{100} and v_{101} which are at distance at least k from each other, at most $k - 1$ from v_1 and v_{10} , and at least k from v_0 and v_{11} . Then, it stops.
 - On TAILS, it produces two calls in cells v_{110} and v_{111} which are at distance at least k from each other, at most $k - 1$ from v_1 and v_{11} , and at least k from v_0 and v_{10} . Then, it stops.

Observe that, the subsequence of the last 4 calls produced by $\mathcal{ADV}\text{-}2$ essentially belongs to the set of sequences of calls produced by $\mathcal{ADV}\text{-}1$.

Now, consider the set of all possible deterministic algorithms \mathcal{A}_2 working on the sequences produced by $\mathcal{ADV}\text{-}2$. 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 $\mathcal{ADV}\text{-}2$ are at distance at least k from 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, i.e., 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 at distance at most $k - 1$ from 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, i.e., 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 $\mathcal{ADV}\text{-}2$ is at least 4. Indeed, in each of the possible sequences

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

generated by the $\mathcal{ADV}\text{-}2$, 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 $\mathcal{ADV}\text{-}2$ is at least 2, which (by Lemma 2) is a lower bound on the competitive ratio of any randomized algorithm for call control.

Next, we show how the adversary $\mathcal{ADV}\text{-}2$ locates the calls in cellular networks of reuse distance $k \geq 2$ completing the proof of the first part of the theorem.

The coordinates of the cells hosting possible calls produced by $\mathcal{ADV}\text{-}2$ are:

$$\begin{aligned} v_0 &= (0, -k, -k) & v_1 &= (0, k, k) \\ v_{00} &= (0, -2k + 1, -2k + 1) & v_{01} &= (0, -1, -1) \\ v_{10} &= (0, 1, 1) & v_{11} &= (0, 2k - 1, 2k - 1) \\ v_{000} &= (k - 1, -2k + 1, -k) & v_{001} &= (-k + 1, -k, -2k + 1) \\ v_{010} &= (k - 1, -k, -1) & v_{011} &= (-k + 1, -1, -k) \\ v_{100} &= (k - 1, 1, k) & v_{101} &= (-k + 1, k, 1) \\ v_{110} &= (k - 1, k, 2k - 1) & v_{111} &= (-k + 1, 2k - 1, k) \end{aligned}$$

An example of all cells hosting possible calls produced by $\mathcal{ADV}\text{-}2$ when $k = 3$ is depicted in Figure 3.




Fig. 3. The calls that may be produced by the adversary $\mathcal{ADV}\text{-}2$ in a cellular network of reuse distance 3. The grey cells host calls of the sequence σ_2^{10} .

We have to show that any of the possible sequences of calls generated according to \mathcal{P}_2 satisfies the constraints defined above. We have four possible sequences to examine: σ_2^{00} , σ_2^{01} , σ_2^{10} and σ_2^{11} . We show that one of them, e.g., σ_2^{10} satisfies the constraints; the proof for the other cases is similar due to symmetry. First, the cells v_0 and v_1 have distance $2k \geq k$. The cells at distance at most $k - 1$ from

v_0 are those contained between the x-rows $k - 1$ and $-k + 1$, between the y-rows $-2k + 1$ and -1 , and between the z-rows $-2k + 1$ and -1 . The cells at distance at most $k - 1$ from v_1 are those contained between the x-rows $k - 1$ and $-k + 1$, between the y-rows 1 and $2k - 1$, and between the z-rows 1 and $2k - 1$. Hence, cells $v_{10}, v_{11}, v_{100}, v_{101}$ are all at distance at least k from v_0 and at most $k - 1$ from v_1 . Since $k \geq 2$, the cells v_{10} and v_{11} have distance $2k - 2 \geq k$. Also, the cells at distance at most $k - 1$ from v_{10} are those contained between the x-rows $k - 1$ and $-k + 1$, between the y-rows $-k + 2$ and k , and between the z-rows $-k + 2$ and k . The cells at distance at most $k - 1$ from v_{11} are those contained between the x-rows $k - 1$ and $-k + 1$, between the y-rows k and $3k - 2$, and between the z-rows k and $3k - 2$. Hence, cells v_{100}, v_{101} are at distance at most $k - 1$ from v_{10} and at least k from v_{11} . In addition, v_{100} and v_{101} are at distance $2k - 2 \geq k$ since $k \geq 2$. \square

4 Conclusions

In this paper, we presented a new online call control algorithm with competitive ratio $16/7$ for cellular networks with reuse distance 2, improving the previous best known upper bound of $7/3$. The algorithm is based on the “classify and randomly select” paradigm, uses only four random bits and works in networks with arbitrarily many frequencies. We have also presented new lower bounds of 2 and 2.5 on the competitiveness against oblivious adversaries of online call control algorithms in cellular networks of reuse distance $k \geq 2$ and $k \geq 6$, respectively. Our new lower bounds improve previous ones for almost all cases of the reuse distance ($k \neq 5$).

An interesting open problem is to close the gap between $16/7$ and 2 on the competitiveness of online randomized call control algorithms in cellular networks with reuse distance 2. In particular, improving the upper bound would require entirely new techniques since the coloring used by algorithm CRS-D seems to be the best possible that satisfies the conditions of Lemma 1.

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 *Proc. 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)*, pp. 531-540, 1996.
4. A. Borodin and R. El-Yaniv. Online Computation and Competitive Analysis. *Cambridge University Press*, 1998.
5. 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 versions in *SPAA '00* and *IPDPS '01*.

6. 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 Online Algorithms (WAOA '03)*, LNCS 2909, Springer, pp. 67-80, 2003.
7. 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.
8. W.K. Hale. Frequency Assignment: Theory and Applications. In *Proc. of the IEEE*, 68(12), pp. 1497-1514, 1980.
9. 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.
10. S. Leonardi, A. Marchetti-Spaccamela, A. Prescutti, and A. Rosen. On-line Randomized Call-Control Revisited. *SIAM Journal on Computing*, Vol. 31(1), pp. 86-112, 2001.
11. R. Motwani and B. Raghavan. Randomized Algorithms. *Cambridge University Press*, 1995.
12. G. Pantziou, G. Pentaris, and P. Spirakis. Competitive Call Control in Mobile Networks. *Theory of Computing Systems*, Vol. 35(6), pp. 625-639, 2002.
13. 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.
14. 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.