Randomized Call Control in Sparse Wireless Cellular Networks*

Ioannis Caragiannis Christos Kaklamanis Evi Papaioannou

Computer Technology Institute and
Department of Computer Engineering and Informatics
University of Patras, 26500 Rio, Greece
E-mail: {caragian,kakl,papaioan}@cti.gr

Abstract

We address an important communication issue arising 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 benefit, i.e., the number of users that communicate without signal interference.

In previous work [4, 5], the authors study the performance of algorithm p–RANDOM; an intuitive on-line randomized call control algorithm for wireless cellular networks where each cell is a regular hexagon. In practice, however, at least parts of the network are expected to be sparser and irregular.

In the present work, we extend previous work on randomized call-control focusing on sparse wireless networks with irregular cell shape and small degree (three or four). We prove analytical upper bounds on the competitive ratio (defined as the maximum over all possible sequences of calls of the ratio of the benefit of the algorithm over the benefit of the optimal off-line algorithm) of algorithm p–RANDOM against oblivious adversaries as a function of the parameter p. Optimizing the upper bound function, we prove that algorithm p–RANDOM is 9/4– and 2.651–competitive on sparse wireless cellular networks supporting one frequency of degree three and four, respectively. Our algorithms extend to networks with many frequencies.

1 Introduction

In the area of mobile communications, which combines wireless and high speed networking technologies, rapid technological progress has been made. It is expected that in the near future, mobile users have access to a wide variety of services available over mobile communication networks.

An architectural approach widely common for mobile networks is the following. A geographical area in which communication takes place is divided into regions (or cells). Each cell is the calling area of a base station. Base stations are connected via a high speed network. In this work, the topology of the high speed network is not of interest. When a mobile user A wishes to communicate with some other user B, a path must be established between the base stations of the cells 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 mobile users are located in the same cell. Networks operating in this way are called wireless cellular networks.

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

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 $B_A(\sigma)$ be the benefit of the on-line algorithm A on the sequence of calls σ, i.e. the number of calls of σ accepted by A and $B_{OPT}(\sigma)$ the benefit of the optimal algorithm OPT. If A is a deterministic algorithm, we define its competitive ratio (or competitiveness) ρ as

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

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

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

where $\mathcal{E}[B_A(\sigma)]$ is the expectation of the number of calls accepted by A, and the maximum is taken 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, 2, 3, 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 planar and arbitrary mobile networks. 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. Lower bounds for call control in arbitrary networks are implied in [3, 12].

The authors in [4] describe algorithm p–RANDOM, an intuitive on-line randomized call control algorithm for networks that support one frequency. Using simple arguments, they prove an upper bound of 2.97 on its competitive ratio against oblivious adversaries in ideal wireless cellular networks. In this way, they beat the barrier of 3 which is a lower bound on the competitiveness of deterministic algorithms in ideal wireless cellular networks. Motivated by algorithm p–RANDOM, they also design another randomized algorithm (algorithm p_δ–RANDOM), for which they prove a competitiveness of 2.934. Also, they present a lower bound of 1.857 on the competitive ratio of any randomized call control algorithm in ideal wireless cellular networks.

In subsequent work [5], using more involved competitive analysis, the authors give stronger bounds on the performance of algorithm p–RANDOM in ideal wireless cellular networks. They present upper and lower bounds of its competitive ratio against oblivious adversaries as a function of the parameter p. Optimizing the upper bound function, they prove that there exists a 2.651–competitive randomized call control algorithm. In this way, they significantly improve the best known upper bound on the competitiveness of on-line randomized call control in ideal wireless cellular networks.

The results in the papers [4, 5] apply only to networks with one frequency but they indicate that randomization helps to beat the deterministic upper bounds.

In this paper, we give new analytical bounds on the competitiveness of algorithm p–RANDOM in sparse wireless cellular networks, i.e., wireless cellular networks in which cells may have irregular shape but the network degree is small. We consider networks of degree three and four; for networks of degree three, we prove that for a specific value of p, algorithm p–RANDOM is 9/4–competitive for networks supporting one frequency and 2.787–competitive for networks supporting many frequencies. We also outline the proof for networks of degree 4; for these networks we can achieve competitive ratios of 2.651 and 3.182 for networks supporting one and many frequencies, respectively. Surprisingly, this matches the best known result for ideal wireless cellular networks presented in [5].

The rest of the paper is structured as follows. In Section 2 we briefly discuss well-known on-line call control algorithms and mention the best known lower bounds on the competitiveness of deterministic and randomized algorithms. In Section 3 we describe in detail algorithm p-RANDOM. In Section 4, we analyze the performance of algorithm p-RANDOM. In Section 4.1, we prove the result for networks of degree four, while in Section 4.2, we outline the proof of the result for networks of degree 4. We conclude in Section 5.

2 Preliminaries

In this section, we briefly describe two well-known on-line algorithms for call control in wireless cellular networks: the greedy algorithm and the CLASSIFY AND RANDOMLY SELECT paradigm. Furthermore, we 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. Also, we discuss lower bounds on the competitiveness of deterministic and randomized on-line call control algorithms, respectively.

Assume that a sequence of calls σ appears in a network that support w frequencies $1, 2, ..., w$. 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, i.e., for the minimum frequency among frequencies $1, 2, ..., w$ 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. [10] have proved that this algorithm is at most $(\Delta + 1)$-competitive against off-line adversaries for networks supporting many frequency (and Δ-competitive for networks supporting one frequency), where Δ is the degree of the network.

The CLASSIFY AND RANDOMLY SELECT paradigm uses a coloring of the cells of the network (coloring of the interference graph) with positive integer (colors) $1, 2, ...$ 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, Awerbuch et al. in [2] (see also Pantziou et al. [10]) prove that the CLASSIFY AND RANDOMLY SELECT algorithm 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. Note that, in wireless cellular networks of degree Δ, the chromatic number of the corresponding interference graph may be up to $\Delta + 1$.

Awerbuch et al. in [1] 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) on-line call control algorithm ALG-1 for one frequency. A call control algorithm ALG for w frequecies 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), i.e.,

```
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.
```

Awerbuch et al. in [1] prove that if ALG–1 is ρ–competitive, then ALG has competitive ratio

$$\frac{1}{1-\exp(-1/\rho)}.$$

In this way, we can prove that the greedy algorithm on networks with many frequencies has slightly better competitive ratio than $\Delta+1$ using the fact that the greedy algorithm for one frequency is Δ–competitive.

Note that Δ is a lower bound on the competitive ratio of any deterministic algorithm. Consider a network with one frequency that consists of a cell v and Δ mutually non–adjacent cells adjacent to v. Now consider the following sequence of calls produced by an adversary which knows how the algorithm makes its decisions. First, a call c appears in cell v. If the algorithm rejects c, the adversary stops the sequence. In this case, the algorithm has no benefit. If the algorithm accepts c, the adversary introduces Δ calls c_1, c_2, ..., c_Δ in cells v_1, v_2, ..., v_Δ, respectively. The algorithm has benefit 1 while the optimal algorithm has benefit Δ by rejecting call c and accepting calls c_1, ..., c_Δ. We conclude that there exist networks of degree Δ in which no deterministic algorithm can be better than Δ–competitive against off–line adversaries.

Recent work of Trevisan [12] on the maximum independent set problem on bounded–degree graphs implies that, in general, the static version of the call control problem on networks of degree Δ is inapproximable within $O\left(\Delta/2^{O(\sqrt{\log\Delta})}\right)$. This means that practical (i.e., algorithms which make their decisions in polynomial time) on-line randomized algorithms with competitive ratio asymptotically better than $O(\Delta^{1-\epsilon})$ for some $\epsilon>0$ are infeasible.

In what follows, we present algorithm p–RANDOM (an intuitive randomized algorithm which performance on ideal wireless cellular networks was analyzed in [4, 5]) and prove that it has significantly better performance than the above mentioned algorithms in sparse wireless cellular networks (i.e., networks of degree three or four). In particular, we prove that its competitive ratio is strictly better than the network degree.

3 Algorithm p–RANDOM

We will describe algorithm p–RANDOM as a call control algorithm for networks supporting one frequency. Then, using the technique described in the previous section, we will obtain a randomized on–line call control algorithm for networks supporting many frequencies. Algorithm p–RANDOM receives as input a sequence of calls in an on–line 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 $p \in [1/\Delta, 1]$, where Δ is the network degree. It is clear that if $p < 1/\Delta$, the competitive ratio is larger than Δ, since the expectation of the benefit of the algorithm on a sequence that consists of only one call is 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 Δ (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 Δ.

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. In this way, we do not have to consider the effects of marking during our analysis.

4 Analysis

In this section we provide the analysis of the competitiveness of algorithm p–RANDOM. We first present the analysis which is not network specific. Then, using our analysis, we prove the upper bound on the competitive ratio of the algorithm for networks of degree three (in Section 4.1) and, we outline the proof of the upper bound for networks of degree four (in Section 4.2).

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 on–line 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 (for some $p \in [0,1]$) on σ. For any call $c \in \sigma$, we denote by $X(c)$ the random variable that indicates whether the algorithm accepted c, i.e.,

$$X(c) = \begin{cases} 0 & \text{if c is rejected} \\ 1 & \text{if c is accepted} \end{cases}$$

Obviously,

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

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

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

Thus, by linearity of expectation,

$$\mathcal{E}[B(\sigma)] = \sum_{c \in A(\sigma)} \left(\mathcal{E} \left[X(c) + \sum_{c' \in \gamma(c)} \frac{X(c')}{d(c')} \right] \right) \tag{1}$$

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

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

Thus, (1) yields

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

$$\geq \quad |A(\sigma)| \min_{c\in A(\sigma)} \left\{ \mathcal{E}\left[X(c) + \sum_{c'\in\gamma'(c)} \frac{X(c')}{d(c')} \right] \right\} \tag{2}$$

In what follows we will try to bound from below the expectation of the random variable

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

for each call $c \in A(\sigma)$.

We concentrate on a call $c \in A(\sigma)$. Let $\Omega = 2^{\gamma'(c)}$ be the set which 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 they appear, they are 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

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

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 of in its effective neighborhood γ. The probability that all calls of γ are rejected given that $\Gamma(c) = \gamma$ is $(1-p)^{|\gamma|}$, while then c is accepted with probability p. Thus,

$$\mathcal{E}[X(c)|\Gamma(c) = \gamma] = p(1-p)^{|\gamma|}. \tag{4}$$

Bounding from below $\mathcal{E}\left[\sum_{c'\in\gamma} \frac{X(c')}{d(c')} |\Gamma(c) = \gamma\right]$ is more complicated. In Sections 4.1 and 4.2, we consider two cases of sparse networks, i.e., networks of degree three and four.

Now, once we have computed a lower bound for $\mathcal{E}[Y(c)]$ for each call $c \in A(\sigma)$, by (2) and the definition of the competitive ratio, we can compute an upper bound on the competitive ratio $\rho(p)$ of algorithm p–RANDOM as

$$\rho(p) \leq \left(\min_{c\in A(\sigma)} \mathcal{E}[Y(c)] \right)^{-1} \tag{5}$$

4.1 Networks of degree three

Assume that the calls of the sequence σ appear in the cells of a network of degree three. Then, for each call c accepted by the optimal algorithm, we consider all possible configurations of the effective neighborhood of c. A configuration is depicted in Figure 3. The call c is the call accepted by the optimal algorithm. This call is adjacent to calls c_1, c_2, c_3. Arrows represent time; for example, the arrow from c_1 to c means that c_1 appeared prior to c. Note that a valid configuration has no directed cycle in the neighborhood of c and arrows between c and its adjacent calls are destined for c.

The sequence σ depicted in Figure 3 also contains the subsequences σ_1 and σ_2. The particular configuration for c will not be considered in our analysis for the following reason. Assume that the optimal algorithm accepts the set of calls $A(\sigma) = \{c\} \cup A(\sigma_1) \cup A(\sigma_2)$, where $A(\sigma_1)$ and $A(\sigma_2)$ denote the subsets of $A(\sigma)$ which consist of calls in σ_1 and σ_2, respectively. Note that $A'(\sigma) = \{c_3\} \cup A(\sigma_1) \cup A(\sigma_2)$

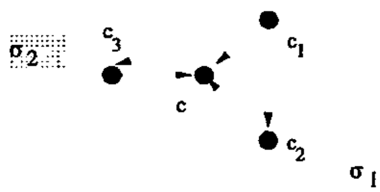

Figure 3: A configuration of the effective neighborhood of an optimal call c.

contains no mutually adjacent calls and, furthermore, has size equal to the size of $A(\sigma)$. Thus, we may assume that $A'(\sigma)$ is the set of optimal calls and consider a much simpler configuration for c_3. So, we obtain another constraint for the configurations we have to consider in our analysis: the optimal call has at most one common neighbor with any of its adjacent calls. In a network with degree three, the six configurations we have to consider are depicted in Figure 4. Symmetric cases have been ommitted.

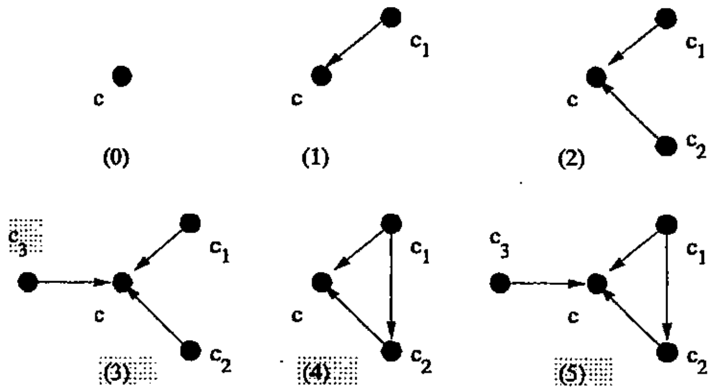

Figure 4: The six configurations we have to consider in our analysis. (i) corresponds to configuration γ_i.

We denote by B_i the expectation of $Y(c)$ given that the effective neighborhood of c has the configuration γ_i, for $i = 0, 1, ..., 5$.

Clearly, by (4), we have that

$$B_0(c) = p \tag{6}$$

For configuration γ_1, we have that $d(c_1) \leq 3$, since the network has maximum degree three. Clearly, $\mathcal{E}\left[\sum_{c' \in \gamma} \frac{X(c')}{d(c')} | \Gamma(c) = \gamma_1\right] \geq p/3$ and, using (4), we obtain that

$$\begin{aligned} B_1(c) &\geq p(1-p) + p/3 \\ &= 4p/3 - p^2 \end{aligned} \tag{7}$$

For configuration γ_2, we again have that $d(c_1) \leq 3$ and $d(c_2) \leq 3$ and $\mathcal{E}\left[\sum_{c' \in \gamma} \frac{X(c')}{d(c')} | \Gamma(c) = \gamma_2\right] \geq 2p/3$. Using (4), we obtain that

$$\begin{aligned} B_2(c) &\geq p(1-p)^2 + 2p/3 \\ &= 5p/3 - 2p^2 + p^3 \end{aligned} \tag{8}$$

For configuration γ_3, we again have that $d(c_1) \leq 3$, $d(c_2) \leq 3$, $d(c_3) \leq 3$ and $\mathcal{E}\left[\sum_{c' \in \gamma} \frac{X(c')}{d(c')} | \Gamma(c) = \gamma_3\right] \geq p$. Using (4), we obtain that

$$B_3(c) \geq p(1-p)^3 + p$$

$$= \quad 2p - 3p^2 + 3p^3 - p^4 \tag{9}$$

For configuration γ_4, we have that $d(c_1) \leq 2$ and $d(c_2) \leq 2$, since the network has maximum degree three. To compute $\mathcal{E}\left[\sum_{c'\in\gamma} \frac{X(c')}{d(c')} | \Gamma(c) = \gamma_4\right]$ observe that c_1 is accepted with probability p while c_2 is accepted with probability p if c_1 was previously rejected, i.e., the probability that c_2 is accepted is $p(1-p)$. Thus,

$$\begin{aligned}\mathcal{E}\left[\sum_{c'\in\gamma} \frac{X(c')}{d(c')} | \Gamma(c) = \gamma_4\right] &\geq p/2 + p(1-p)/2 \\ &= p - p^2/2\end{aligned}$$

Using (4), we obtain that

$$\begin{aligned}B_4(c) &\geq p(1-p)^2 + p - p^2/2 \\ &= 2p - 5p^2/2 + p^3 \end{aligned}\tag{10}$$

For configuration γ_5, we have that $d(c_1) \leq 2$, $d(c_2) \leq 2$ and $d(c_3) \leq 3$, since the network has maximum degree three. Using the same reasoning as above, we have

$$\begin{aligned}\mathcal{E}\left[\sum_{c'\in\gamma} \frac{X(c')}{d(c')} | \Gamma(c) = \gamma_5\right] &\geq p/2 + p(1-p)/2 + p/3 \\ &= 4p/3 - p^2/2\end{aligned}$$

Using (4), we obtain that

$$\begin{aligned}B_4(c) &\geq p(1-p)^3 + 4p/3 - p^2/2 \\ &= 7p/3 - 7p^2/2 + 3p^3 - p^4 \end{aligned}\tag{11}$$

Now, the expectation of $Y(c)$ can be expressed as

$$\mathcal{E}[Y(c)] \geq \min_i B_i(c) \tag{12}$$

By making simple calculations with (7)–(11), we may verify that (since $p \in [0, 1]$)

$$B_1(c) \leq \min\{B_2(c), B_3(c), B_4(c), B_5(c)\}.$$

For $p \geq 1/3$ (which contains the range of p in which we are interested for improving known results), (6) and (7) also yield $B_1(c) \leq B_0(c)$. Thus, for $p \in [1/3, 1]$, $B_1(c)$ is a lower bound for $\mathcal{E}[Y(c)]$ for each call $c \in A(\sigma)$ which only depends on p. Using (5), we obtain that for $p \in [1/3, 1]$, the competitive ratio of algorithm p–RANDOM is

$$\rho(p) \leq \frac{3}{4p - 3p^2}.$$

The right side of the above inequality is minimized for $p = 2/3$ to $9/4$. We have obtained the following theorem.

Theorem 1 *There exists a 9/4–competitive randomized on-line call control algorithm for sparse wireless cellular networks of degree three that support one frequency.*

Using the technique of Awerbuch et al., we can transform this algorithm to work with many frequencies. We obtain the following corollary.

Corollary 2 *There exists a 2.787–competitive randomized on-line call control algorithm for sparse wireless cellular networks of degree three that support many frequencies.*

4.2 Networks of degree four

In this section we outline the proof for the competitiveness of algorithm p–RANDOM on networks of degree four. Assume that the calls of the sequence σ appear in the cells of a network of degree four. Again, for each call c accepted by the optimal algorithm, we consider all possible configurations of the effective neighborhood of c. These configuration are constructed as in the previous section and now have the following constraints:

- At most four calls appear prior to c.
- There is no directed cycle in the neighborhood of c and arrows between c and its adjacent calls are destined for c.
- The optimal call c has at most two common neighbors with any of its adjacent calls.

In a network with degree four, we have to consider twenty three configurations instead of six that we considered for networks of degree three. Due to the limited space, we will not list all these configurations here. We denote by $B'_i(c)$ the expectation of $Y(c)$ given that the effective neighborhood of c has the configuration γ'_i, for $i = 0, 1, ..., 22$. We will just compute the expectation $B'_4(c)$ of $Y(c)$ given that the effective neighborhood of c has the configuration γ'_4 which is the same with configuration γ_4 of Figure 4.

We have that $d(c_1) \leq 3$ and $d(c_2) \leq 3$, since the network has maximum degree four. To compute $\mathcal{E}\left[\sum_{c'\in\gamma} \frac{X(c')}{d(c')}|\Gamma(c) = \gamma'_4\right]$, observe that c_1 is accepted with probability p while c_2 is accepted with probability p if c_1 was previously rejected, i.e., the probability that c_2 is accepted is $p(1-p)$. Thus,

$$\begin{aligned}\mathcal{E}\left[\sum_{c'\in\gamma} \frac{X(c')}{d(c')}|\Gamma(c) = \gamma_4\right] &\geq p/3 + p(1-p)/3 \\ &= 2p/3 - p^2/3\end{aligned}$$

Using (4), we obtain that

$$\begin{aligned}B'_4(c) &\geq p(1-p)^2 + 2p/3 - p^2/3 \\ &= 5p/3 - 7p^2/3 + p^3 \qquad (13)\end{aligned}$$

Using similar reasoning as in Section 4.1, by examining all configurations, we can verify that, for some range of p, $B'_4(c) = \min_i B'_i(c)$. Thus, for this specific range of p, $B'_4(c)$ is a lower bound for $\mathcal{E}[Y(c)]$ for each call $c \in A(\sigma)$ which only depends on p. Using (5), we obtain that for the specific range of p, the competitive ratio $\rho'(p)$ of algorithm p–RANDOM is

$$\rho'(p) \leq \frac{3}{5p - 7p^2 + 3p^3}.$$

The right side of the above inequality is minimized for $p = 5/9$ to $729/275 = 2.651$. We have obtained the following theorem.

Theorem 3 *There exists a 2.651–competitive randomized on–line call control algorithm for sparse wireless cellular networks of degree four that support one frequency.*

Surprisingly, this bound matches the best known result for ideal wireless cellular networks presented in [5]. Using the technique of Awerbuch et al., we can transform this algorithm to work with many frequencies. We obtain the following corollary.

Corollary 4 *There exists a 3.182–competitive randomized on–line call control algorithm for sparse wireless cellular networks of degree four that support many frequencies.*

5 Conclusion

We have presented new analytical bounds on the competitiveness of algorithm p–RANDOM in sparse wireless cellular networks, i.e., wireless cellular networks in which cells may have irregular shape but the network degree is small (three or four). Our analysis is nearly–tight; the interested reader may verify that there exist networks of degree three and four supporting one frequency such that algorithm p–RANDOM is at least 9/4 and 2.56, respectively. An interesting open problem is to improve the upper bounds for networks supporting many frequencies. Also, designing better randomized algorithms is another interesting open problem.

References

[1] B. Awerbuch, Y. Azar, A. Fiat, S. Leonardi, and A. Rosen. Competitive On–line Call Admission in Optical Networks. In *Proceedings of the 4th Annual European Symposium on Algorithms (ESA '96)*, LNCS 1136, Springer, pp. 431–444, 1996.

[2] B. Awerbuch, Y. Bartal, A. Fiat, A. Rosen. Competitive Non–Preemptive Call Control. In *Proceedings 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 Communication in Cellular Networks. In *Proc. of the 12th Annual ACM Symposium on Parallel Algorithms and Architectures (SPAA '00)*, pp. 46–53, 2000.

[5] I. Caragiannis, C. Kaklamanis, and E. Papaioannou. Competitive Analysis of On–line Randomized Call Control in Cellular Networks. In *Proc. of the 15th International Parallel and Distributed Processing Symposium (IPDPS '01)*, 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. In *Proceedings of 15th Annual Symposium on Theoretical Aspects of Computer Science (STACS '98)*, LNCS 1373, Springer, pp. 3–13, 1998.

[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. In *Proceedings of International Symposium on Algorithms and Computation (ISAAC '97)*, LNCS 1350, Springer, pp. 404–413, 1997.

[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] L. Trevisan. Non–Approximability Results for Optimization Problems on Bounded Degree Instances. In *Proc. of the 33rd Annual ACM Symposium on Theory of Computing (STOC '01)*, pp. 453–461, 2001.