Alcom-FT Technical Report Series
ALCOMFT-TR-01-60

Efficient On–line Communication in Cellular Networks *

Ioannis Caragiannis † Christos Kaklamanis ‡ Evi Papaioannou §

ABSTRACT

In this paper we consider communication issues arising in mobile networks that utilize Frequency Division Multiplexing (FDM) technology. In such networks, many users within the same geographical region can communicate simultaneously with other users of the network using distinct frequencies. The spectrum of available frequencies is limited; thus, efficient solutions to the frequency allocation and the call control problem are essential. In the frequency allocation problem, given users that wish to communicate, the objective is to minimize the required spectrum of frequencies so that communication can be established without signal interference. 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 served. We consider cellular, planar, and arbitrary network topologies.

In particular, we study the on–line version of both problems using competitive analysis. For frequency allocation in cellular networks, we improve the best known competitive ratio upper bound of 3 achieved by the folklore Fixed Allocation algorithm, by presenting an almost tight competitive analysis for the greedy algorithm; we prove that its competitive ratio is between 2.429 and 2.5. For the call control problem, we present the first randomized algorithm that beats the deterministic lower bound of 3 achieving a competitive ratio of 2.934 in cellular networks. Our analysis has interesting extensions to arbitrary networks. Also, using Yao's Minimax Principle, we prove two lower bounds of 1.857 and 2.086 on the competitive ratio of randomized call control algorithms for cellular and arbitrary planar networks, respectively.

1. INTRODUCTION

In the area of mobile computing, 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 communication networks.

An architectural approach widely common for wireless networks 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. 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 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 A and B are located in the same region.

Many users of the same region can communicate simultaneously with other users of the network. 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 both within the same region and adjacent regions. Since the spectrum of available frequencies is limited, important engineering problems related to the efficient use of the frequency spectrum arise.

The network topology usually adopted is a finite portion of the infinite triangular lattice. This results from the uniform distribution of 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. Associated with the cellular network is an *in-*

*This work was partially funded by the European Union (IST FET Project ALCOM–FT and RTN Project ARACNE) and the Greek General Secretariat of Research and Technology (Programme ΠΕΝΕΔ 99).

†Computer Technology Institute and Department of Computer Engineering and Informatics, University of Patras, 26500 Rio, Greece. E–mail: caragian@cti.gr

‡Computer Technology Institute and Department of Computer Engineering and Informatics, University of Patras, 26500 Rio, Greece. E–mail: kakl@cti.gr

§Computer Technology Institute and Department of Computer Engineering and Informatics, University of Patras, 26500 Rio, Greece. E–mail: papaioan@cti.gr

terference graph G that reflects possible signal interference. Vertices correspond to cells and an edge (u, v) exists in the graph if and only if the cells corresponding to u and v are adjacent. Due to geometry, we call this graph a hexagon graph (see Figure 1). If the above assumptions (uniform distribution of base stations and equivalence of transmitters) do not hold, arbitrary interference graphs can be used to model the underlying network. At the rest of this paper, we use the term cellular network especially for networks with hexagon interference graph, like the one depicted in Figure 1. We use the terms planar and arbitrary network for networks with planar and arbitrary interference graphs, respectively.

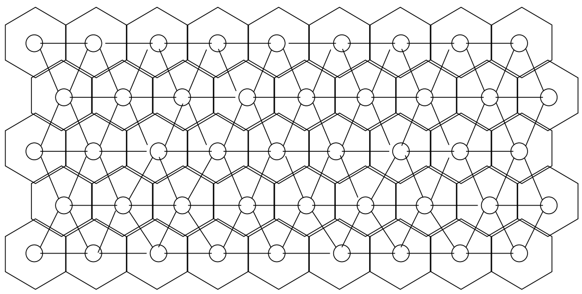

Figure 1: A cellular network and the corresponding interference (hexagon) graph

Given users that wish to communicate, two problems are of main interest:

- The *frequency allocation* problem is to assign frequencies to the users so that signal interference is avoided, minimizing the total number of frequencies used.
- The *call control* (or call admission) problem on a network that supports a spectrum of w available frequencies is to assign frequencies to users so that signal interference is avoided, maximizing the number of users served.

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, an on–line frequency allocation algorithm accepts the call assigning a frequency to it, while 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 A_{OPT}.

For frequency allocation algorithms, let $C_A(\sigma)$ be the cost of the on–line algorithm A on the sequence of calls σ, i.e. the number of frequencies used by A and $C_{OPT}(\sigma)$ the cost of the optimal algorithm A_{OPT}. If A is a deterministic algorithm, we define its competitive ratio ρ as

$$\rho = \max_{\sigma} \frac{C_A(\sigma)}{C_{OPT}(\sigma)},$$

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

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

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

For call control algorithms, 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 A_{OPT}. If A is a deterministic algorithm, we define its competitive ratio 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.

Frequency allocation in cellular networks has recently received much attention. The static version of the problem is studied in [9] and [7], while [5] deals with the on–line version of the problem. Janssen et al. [5] use competitive analysis to evaluate the performance of several frequency allocation algorithms and prove lower bounds. Among other results, they prove that no on–line deterministic frequency allocation algorithm can achieve a competitive ratio better than 2 while they mention the folklore Fixed Allocation algorithm which is 3–competitive.

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 recently studied in [10], [2], and [3]. Pantziou et al. [10] present upper bounds for planar and arbitrary networks. Lower bounds for call control in arbitrary networks are presented in [2]. The authors in [3] study intuitive call control algorithms in cellular networks and prove a lower bound on the performance of randomized call control algorithms.

In this paper, we address the on–line version of the frequency allocation and the call control problem. For frequency allocation in cellular networks, we improve the best known competitive ratio upper bound of 3 achieved by the Fixed Allocation algorithm [5] presenting an almost tight competitive analysis for the greedy frequency allocation algorithm (Section 2). In particular, we prove that its competitive ratio is between 2.429 and 2.5. For the call control problem, we present the first randomized algorithm that beats the deterministic lower bound of 3 achieving a competitive ratio of 2.934 on cellular networks (Section 3.1). This upper bound holds for networks that support one frequency, but gives strong evidence that using randomization the known bounds can be improved. Our analysis has interesting extensions to arbitrary networks that support many frequencies. Also, using Yao's Minimax Principle, we prove two lower bounds of 1.857 and 2.086 on the competitive ratio of randomized call control algorithms for cellular and arbitrary planar networks, respectively (Section 3.2). Our algorithms for both problems are simple and can be easily implemented with small communication overhead (exchange of messages) between the base stations of the network.

2. GREEDY FREQUENCY ALLOCATION

An intuitive deterministic algorithm for the on–line frequency allocation problem is the greedy algorithm which is the following. The algorithm considers the frequencies as positive integers $1, 2, \ldots$. When a new call appears, it is assigned the minimum available frequency so that the call does not interfere (is not assigned the same frequency) with calls of the same or adjacent cells. In this section we present an almost tight competitive analysis of the greedy algorithm.

THEOREM 1. *The greedy frequency allocation algorithm is at most 2.5–competitive against an off–line adversary when applied to a cellular network.*

PROOF. Let σ be a sequence of calls in a cellular network and D be the maximum number of calls in any three mutually adjacent cells of the network. Obviously, D is a lower bound on the number of frequencies required for an optimal frequency allocation to σ.

Consider the execution of the greedy algorithm on σ, and let c_0 be the cell that contains the call that was assigned the highest frequency a_0. We will prove that $a_0 \leq 2.5D$, proving that the greedy algorithm is at most 2.5–competitive against an off–line adversary.

We denote by x_0 the number of calls in cell c_0. By the definition of the greedy algorithm, since the frequency a_0 has been assigned to a call of cell c_0, the frequencies $1, \ldots, a_0 - 1$ must also have been assigned to calls of c_0 and its surrounding cells. Note that the number of all these calls is at most $3D - 2x_0$. Thus, we obtain the following constraint for x_0.

$$x_0 \leq \frac{3D - a_0}{2} \tag{1}$$

We call the six surrounding cells of c_0, $c_1, c_2, c_3, c_4, c_5, c_6$, so that for the corresponding highest frequency a_1, a_2, a_3, a_4, a_5, a_6 assigned to calls of these cells, the following inequality holds:

$$a_0 > a_1 \geq a_2 \geq a_3 \geq a_4 \geq a_5 \geq a_6$$

We also denote by $y_i(j)$, for $0 \leq i \neq j \leq 6$, the number of calls in cell c_j that have been assigned higher frequencies than a_i. Obviously, it is $y_i(j) \leq x_j$, for $0 \leq i \neq j \leq 6$.

Consider frequency a_1 assigned to some call of cell c_1. It is $a_1 = a_0 - y_1(0)$. Furthermore, a_1 is upper–bounded by the number of calls in c_1 and its surrounding cells, ignoring the $y_1(0)$ calls of c_0 (which have been assigned frequencies higher than a_1). We obtain that

$$x_1 \leq \frac{3D - a_0}{2} \tag{2}$$

For the frequency a_2 assigned to some call of c_2, it is

$$a_2 \geq a_0 - y_2(0) - y_2(1) \tag{3}$$

We now distinguish between two main cases.

CASE I: The cell c_2 is adjacent to c_1.

In this case, a_2 is upper–bounded by the number of calls in c_2 and its surrounding cells, ignoring the $y_2(0)$ calls of c_0 and the $y_2(1)$ calls of c_1 (which have been assigned frequencies higher than a_2). The number of calls in c_2 and its surrounding cells is at most $2D - x_2 + x_0 + x_1$. Using (3), we obtain that

$$x_2 \leq 2D + x_0 + x_1 - a_0 \tag{4}$$

Note that the number of calls in c_0 and its surrounding cells, which is an upper bound for a_0, is at most $2D - x_0 + x_1 + x_2$. Using (2) and (4), we obtain that

$$\begin{aligned} a_0 &\leq 2D - x_0 + x_1 + x_2 \\ &\leq 4D + 2x_1 - a_0 \\ &\leq 7D - 2a_0 \Rightarrow \\ a_0 &\leq \frac{7D}{3}. \end{aligned}$$

CASE II: The cell c_2 is not adjacent to c_1.

In this case, a_2 is upper–bounded by the number of calls in c_2 and its surrounding cells, ignoring the $y_2(0)$ calls of c_0 (which have been assigned frequencies higher than a_2). The number of calls in c_2 and its surrounding cells is at most $3D - 2x_2$. Using (3), we obtain that

$$x_2 \leq \frac{3D - a_0 + y_2(1)}{2} \tag{5}$$

For the frequency a_3 assigned to some call of c_3, it is

$$a_3 \geq a_0 - y_3(0) - y_3(1) - y_3(2) \tag{6}$$

We now distinguish between the following three cases.

CASE II.1: The cell c_3 is adjacent to c_1 (and possibly to c_2).

In this case, a_3 is upper–bounded by the number of calls in c_3 and its surrounding cells, ignoring the $y_3(0)$ calls of

c_0 and the $y_3(1)$ calls of c_1 (which have been assigned frequencies higher than a_3). The number of calls in c_3 and its surrounding cells is at most $2D - x_3 + x_0 + x_1$. Using (6), we obtain that

$$x_3 \leq 2D + x_0 + x_1 + y_3(2) - a_0 \tag{7}$$

Note that the number of calls in c_0 and its surrounding cells, which is an upper bound for a_0, is at most $2D - x_0 + x_1 + x_3$. Using (2), (5), (7), and the fact that $y_3(2) \leq x_2$ and $y_2(1) \leq x_1$, we obtain that

$$\begin{aligned}
a_0 &\leq 2D - x_0 + x_1 + x_3 \\
&\leq 4D + 2x_1 + y_3(2) - a_0 \\
&\leq 4D + 2x_1 + x_2 - a_0 \\
&\leq \frac{11D}{2} + 2x_1 + \frac{y_2(1)}{2} - \frac{3a_0}{2} \\
&\leq \frac{11D}{2} + \frac{5x_1}{2} - \frac{3a_0}{2} \\
&\leq \frac{37D}{4} - \frac{11a_0}{4} \Rightarrow \\
a_0 &\leq \frac{37D}{15}.
\end{aligned}$$

CASE II.2: The cell c_3 is adjacent to c_2 (but not to c_1).

In this case, a_3 is upper–bounded by the number of calls in c_3 and its surrounding cells, ignoring the $y_3(0)$ calls of c_0 and the $y_3(2)$ calls of c_2 (which have been assigned frequencies higher than a_3). The number of calls in c_3 and its surrounding cells is at most $2D - x_3 + x_0 + x_2$. Using (6), we obtain that

$$x_3 \leq 2D + x_0 + x_2 + y_3(1) - a_0 \tag{8}$$

Note that the number of calls in c_0 and its surrounding cells, which is an upper bound for a_0, is at most $2D - x_0 + x_2 + x_3$. Using (2), (5), (8), and the fact that $y_3(1) \leq x_1$ and $y_2(1) \leq x_1$, we obtain that

$$\begin{aligned}
a_0 &\leq 2D - x_0 + x_2 + x_3 \\
&\leq 4D + 2x_2 + y_3(1) - a_0 \\
&\leq 4D + x_1 + 2x_2 - a_0 \\
&\leq 7D + x_1 + y_2(1) - 2a_0 \\
&\leq 7D + 2x_1 - 2a_0 \\
&\leq 10D - 3a_0 \Rightarrow \\
a_0 &\leq \frac{5D}{2}.
\end{aligned}$$

CASE II.3: The cell c_3 is not adjacent to c_1 and c_2.

In this case, a_3 is upper–bounded by the number of calls in c_3 and its surrounding cells, ignoring the $y_3(0)$ calls of c_0 (which have been assigned frequencies higher than a_3). The number of calls in c_3 and its surrounding cells is at most $3D - 2x_3$. Using (6), we obtain that

$$x_3 \leq \frac{3D - a_0 + y_3(1) + y_3(2)}{2} \tag{9}$$

For the frequency a_4 assigned to some call of c_4, it is

$$a_4 \geq a_0 - y_4(0) - y_4(1) - y_4(2) - y_4(3) \tag{10}$$

Again, we distinguish between the following three subcases.

CASE II.3.a: The cell c_4 is adjacent to c_1 and c_2.

In this case, a_4 is upper–bounded by the number of calls in c_4 and its surrounding cells, ignoring the $y_4(0)$ calls of c_0, the $y_4(1)$ calls of c_1, and the $y_4(2)$ calls of c_2 (which have been assigned frequencies higher than a_4). The number of calls in c_4 and its surrounding cells is at most $2D - x_4 + x_0 + x_1$. Using (10), we obtain that

$$x_4 \leq 2D + x_0 + x_1 + y_4(3) - a_0 \tag{11}$$

Note that the number of calls in cells c_5 and c_6 which are adjacent to c_3 is $x_5, x_6 \leq D - x_0 - x_3$. Thus, the number of calls in c_0 and its surrounding cells, which is an upper bound for a_0, is at most $2D - x_0 - x_3 + x_1 + x_2 + x_4$. Using (2), (5), (9), (11), and the fact that $y_4(3) \leq x_3$ and $y_2(1) \leq x_1$, we obtain that

$$\begin{aligned}
a_0 &\leq 2D - x_0 - x_3 + x_1 + x_2 + x_4 \\
&\leq 4D - x_3 + 2x_1 + x_2 + y_4(3) - a_0 \\
&\leq \frac{11D}{2} + 2x_1 + \frac{y_2(1)}{2} - \frac{3a_0}{2} \\
&\leq \frac{11D}{2} + \frac{5x_1}{2} - \frac{3a_0}{2} \\
&\leq \frac{37D}{4} - \frac{11a_0}{4} \Rightarrow \\
a_0 &\leq \frac{37D}{15}.
\end{aligned}$$

CASE II.3.b: The cell c_4 is adjacent to c_1 and c_3.

In this case, a_4 is upper–bounded by the number of calls in c_4 and its surrounding cells, ignoring the $y_4(0)$ calls of c_0, the $y_4(1)$ calls of c_1, and the $y_4(3)$ calls of c_3 (which have been assigned frequencies higher than a_4). The number of calls in c_4 and its surrounding cells is at most $2D - x_4 + x_0 + x_1$. Using (10), we obtain that

$$x_4 \leq 2D + x_0 + x_1 + y_4(2) - a_0 \tag{12}$$

Note that the number of calls in c_0 and its surrounding cells, which is an upper bound for a_0, is at most $2D - x_0 + x_1 + x_4$. Using (2), (5), (9), (12), and the fact that $y_4(2) \leq x_2$ and $y_2(1) \leq x_1$, we obtain that

$$\begin{aligned}
a_0 &\leq 2D - x_0 + x_1 + x_4 \\
&\leq 4D + 2x_1 + y_4(2) - a_0 \\
&\leq 4D + 2x_1 + x_2 - a_0 \\
&\leq \frac{11D}{2} + 2x_1 + \frac{y_2(1)}{2} - \frac{3a_0}{2} \\
&\leq \frac{11D}{2} + \frac{5x_1}{2} - \frac{3a_0}{2} \\
&\leq \frac{37D}{4} - \frac{11a_0}{4} \Rightarrow \\
a_0 &\leq \frac{37D}{15}.
\end{aligned}$$

CASE II.3.c: The cell c_4 is adjacent to c_2 and c_3.

In this case, a_4 is upper–bounded by the number of calls in c_4 and its surrounding cells, ignoring the $y_4(0)$ calls of c_0, the $y_4(2)$ calls of c_2, and the $y_4(3)$ calls of c_3 (which have been assigned frequencies higher than a_4). The number of calls in c_4 and its surrounding cells is at most $2D - x_4 + x_0 + x_2$. Using (10), we obtain that

$$x_4 \leq 2D + x_0 + x_2 + y_4(1) - a_0 \tag{13}$$

Note that the number of calls in c_0 and its surrounding cells, which is an upper bound for a_0, is at most $2D - x_0 + x_2 + x_4$. Using (2), (5), (9), (13), and the fact that $y_4(1) \leq x_1$ and $y_2(1) \leq x_1$, we obtain that

$$\begin{aligned}
a_0 &\leq 2D - x_0 + x_2 + x_4 \\
&\leq 4D + 2x_2 + y_4(1) - a_0 \\
&\leq 4D + x_1 + 2x_2 - a_0 \\
&\leq 7D + x_1 + y_2(1) - 2a_0 \\
&\leq 7D + 2x_1 - 2a_0 \\
&\leq 10D - 3a_0 \Rightarrow \\
a_0 &\leq \frac{5D}{2}.
\end{aligned}$$

In any case, it is $a_0 \leq 2.5D$. The theorem follows. $\square$

We now give a lower bound on the performance of the greedy algorithm.

THEOREM 2. *The greedy algorithm is at least* 2.429*–competitive against an off–line adversary when applied to a cellular network.*

PROOF. Consider the cellular network and the sequence of calls σ' shown in the upper part of Figure 2. Calls that appear in step 1 are assigned frequency 1. At any subsequent step $2 \leq i \leq 17$, the greedy algorithm will assign to all the calls that appear in step i the frequency i, since the frequencies $1, 2, ..., i-1$ are already used by calls in the same or adjacent cells.

In this counterexample, the greedy algorithm uses 17 frequencies while an optimal frequency allocation of σ' with 7 frequencies is depicted in the lower part of Figure 2. Thus, the competitive ratio of the greedy algorithm is

$$\rho = \max_{\sigma} \frac{C_A(\sigma)}{C_{OPT}(\sigma)} \geq \frac{C_A(\sigma')}{C_{OPT}(\sigma')} = \frac{17}{7} = 2.429.$$

$\square$

3. RANDOMIZED CALL CONTROL

3.1 Upper bounds

In this section we prove that, at least for networks that support one frequency, the use of randomization yields on–line call control algorithms with competitive ratios better than the lower bound for deterministic ones.

We consider the algorithm p–RANDOM described as follows.

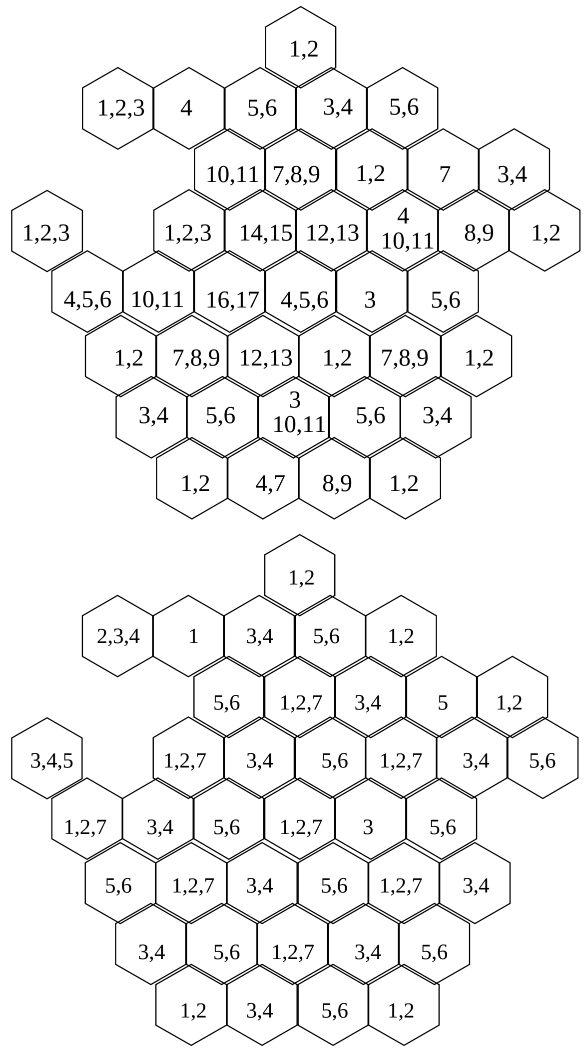

Figure 2: The lower bound on the performance of the greedy algorithm. In the upper part integers correspond to the step in which a call appears. An optimal allocation of frequencies is depicted in the lower part.

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 c is in cell 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.

Marking cells on rejection guarantees that the 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 c and one call in 3 (mutually not adjacent) cells adjacent to c. The probability that the randomized algorithm does not accept a call in cell c drops exponentially as t increases, and the benefit approaches 1, while the optimal benefit is 3.

We prove the following.

THEOREM 3. *For $1/3 < p < 1$, the p–*RANDOM *algorithm has competitive ratio at most*

$$\frac{3}{1+(1-p)^6(3p-1)}$$

against oblivious adversaries.

PROOF. Let σ be a sequence of calls and consider the execution of p–RANDOM on σ. For any call $c \in \sigma$, we denote by $X(c)$ the random variable that indicates whether the algorithm accepted c. 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, and $d(c)$ is the number of calls in $A(\sigma)$ that are in cells adjacent to the cell of c. It is

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

Let q be the probability that, when the call c is presented to the algorithm, no call has been accepted in surrounding cells. Then, $\mathcal{E}[\bar{b}](c) \geq pq + \frac{1-q}{3}$. Note that for any sequence of calls, it is $q \geq (1-p)^6$. Thus, we obtain that

$$\mathcal{E}[\bar{b}](c) \geq (1-p)^6 \left(p - \frac{1}{3}\right) + \frac{1}{3}.$$

The theorem follows by linearity of expectations. $\square$

The expression in Theorem 3 is minimized to 2.97 for $p = 3/7$. We can slightly improve this bound by using different probabilities p_δ for the random choice in lines 6 and 7 of the algorithm according to the number δ of surrounding cells of v in which calls have already been presented. Note that, when considering a call c in cell v, if calls have been presented in all the six cells adjacent to v and the algorithm enters lines 6–7, there is no need to reject c. So, we can use $p_6 = 1$. By setting appropriately the values of the probabilities ($p_0 = p_1 = p_2 = p_3 = 0.377$, $p_4 = 0.472$, and $p_5 = 0.497$) and following more complicated analysis, we can prove the next theorem. The details are omitted.

THEOREM 4. *There exists a randomized call control algorithm for cellular networks that support one frequency with competitive ratio 2.934 against oblivious adversaries.*

We, now, extend our analysis for arbitrary networks. We express the result in terms of the maximum degree Δ of the network. Pantziou et al. [10] study the competitiveness of the $\frac{1}{\Delta}$–RANDOM algorithm on the average case, and prove that its competitive ratio on sequences of calls generated according to specific probability distributions is significantly smaller than the average degree of the network.

THEOREM 5. *For any network with maximum degree $\Delta \geq 2$, there exists a 0.985Δ–competitive randomized call control algorithm against oblivious adversaries.*

PROOF. Let G be a network with maximum degree Δ. Consider the execution of algorithm p–RANDOM on a sequence σ of calls on G. Following the notation of Theorem 3, we have that the expectation of the amortized benefit for each call c accepted by the optimal algorithm is

$$\mathcal{E}[\bar{b}](c) \geq (1-p)^\Delta \left(p - \frac{1}{\Delta}\right) + \frac{1}{\Delta}.$$

The right part of the expression maximized for $p = \frac{2}{\Delta+1}$ to

$$\left(1 - \frac{2}{\Delta+1}\right)^\Delta \left(\frac{2}{\Delta+1} - \frac{1}{\Delta}\right) + \frac{1}{\Delta} \geq \frac{e^{-2}/9+1}{\Delta},$$

where the inequality is obtained by straightforward calculations using the fact that $(1-1/n)^{n-1} \geq e^{-1}$, for $n > 1$.

We conclude that, for the benefit of the algorithm $\frac{2}{\Delta+1}$–RANDOM, it holds

$$B_{OPT}(\sigma) \leq \frac{\Delta}{e^{-2}/9+1}\mathcal{E}[B(\sigma)] = 0.985\Delta\mathcal{E}[B(\sigma)].$$

The theorem follows. $\square$

Note that using a theorem of Awerbuch et al. [1], theorem 5 implies the existence of a randomized call control algorithm for arbitrarily many frequencies with competitive ratio $0.985\Delta + 1$.

Consider a network G and let $\Gamma(G)$ be the maximum independent set in the neighbourhood of each node. For any network, $\Gamma(G)$ is a lower bound on the competitive ratio of any deterministic algorithm [3]. By inserting $\Gamma(G)$ in the proof of Theorem 5, we obtain the following.

COROLLARY 6. *For any network that supports one frequency with $\Gamma(G) \geq 2$ and $1/2 < p < 1$, the algorithm p–RANDOM has (strictly) better competitive ratio than any deterministic algorithm.*

Obviously, if $\Gamma(G) = 1$ or $\Gamma(G) = 0$, the greedy deterministic algorithm is optimal.

3.2 Lower bounds

Using the minimax principle proposed by Yao [12] (see also [8]), we prove two lower bounds on the competitive ratio against oblivious adversaries of any randomized algorithm on cellular and arbitrary planar networks. We consider networks that support one frequency; our lower bounds can be trivially extended for networks that support multiple frequencies. In our proofs, we use the following lemma.

LEMMA 7 (MINIMAX PRINCIPLE [8]). *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 for call control in cellular networks is the following.

THEOREM 8. *No randomized call–control algorithm can be better than 1.857–competitive against an oblivious adversary when applied to a cellular network.*

PROOF. Consider a proper coloring of the cells of the network with the colors RED, BLUE, and GREEN. Let r_0 be a red cell, b_1, b_2, and b_3 its blue neighbors, and g_1, g_2, and g_3 its green neighbors. We will prove that there exists an adversary $\mathcal{ADV}$ that produces calls according to a probability distribution $\mathcal{P}$ in such way that no deterministic algorithm can be better than 1.857–competitive under $\mathcal{P}$ even if it knows the probability distribution $\mathcal{P}$ in advance.

We define the probability distribution $\mathcal{P}$ as follows. First, the adversary produces a call in the red cell r_0. Then, it

- either stops, with probability $4/7$,
- or does the following, with probability $3/7$. It presents two calls in the cells b_1 and b_2, and
  - either produces a call in the cell b_3 with probability $1/3$,
  - or presents three calls in the cells g_1, g_2, and g_3, with probability $2/3$.

It can be easily seen that the expected benefit of the optimal off–line algorithm on sequences of calls generated according to $\mathcal{P}$ is

$$\mathcal{E}_{\mathcal{P}}[B_{OPT}(\sigma)] = 1 \cdot \frac{4}{7} + 3 \cdot \frac{3}{7} = \frac{13}{7}.$$

Let A be a deterministic call control algorithm that runs on the calls produced by $\mathcal{ADV}$. 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 r_0, and q_1 the number of executions in which A accepts both calls in cells b_1 and b_2.

The expected number of executions in which the algorithm does not accept the call in cell r_0 and the adversary produces a call cell b_3 is $\frac{3}{7}\frac{1}{3}(t-q_0)$. Similarly, the expected number of executions in which the algorithm does not accept the calls in cells r_0, b_1, and b_2 and the adversary produces calls in cells g_1, g_2, g_3 is $\frac{2}{3}\left(\frac{3}{7}(t-q_0)-q_1\right)$. Thus,

$$\begin{aligned}\mathcal{E}_{\mathcal{P}}[B_A(\sigma)] &\leq \frac{q_0 + 2q_1 + \frac{3}{7}\frac{1}{3}(t-q_0) + 3\frac{2}{3}\left(\frac{3}{7}(t-q_0)-q_1\right)}{t} \\ &= 1\end{aligned}$$

and

$$c_A^{\mathcal{P}} \geq \frac{13}{7} = 1.857.$$

By Lemma 7, we obtain that this is a lower bound on the competitive ratio of any randomized algorithm against an oblivious adversary. □

The lower bound for planar networks follows. Note that the best known upper bound is 5 achieved by the "classify and randomly select" algorithm [1, 10].

THEOREM 9. *There exists a planar network on which no randomized call–control algorithm can be better than 2.086–competitive against oblivious adversaries.*

PROOF. Consider a planar network with the interference graph shown in Figure 3. We will prove that there exists an adversary $\mathcal{ADV}$ that produces calls according to a probability distribution $\mathcal{P}$ in such way that no deterministic algorithm can be better than 2.086–competitive under $\mathcal{P}$ even if it knows the probability distribution $\mathcal{P}$ in advance.

We define the probability distribution $\mathcal{P}$ as follows. First, the adversary produces a call in cell c_0. Then, it

- either stops, with probability $4/5$,
- or does the following, with probability $1/5$. It presents a call in cells $c_1, ..., c_5$, and
  - either stops with probability $2/7$,
  - or presents a call in cells $c_6, ..., c_{12}$, with probability $5/7$.

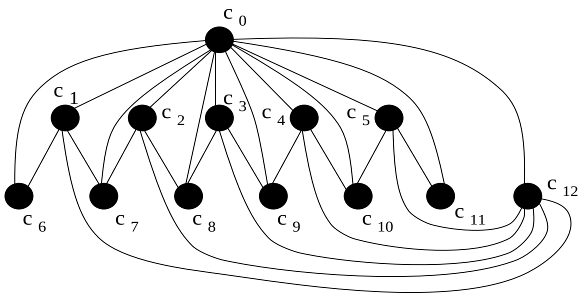

Figure 3: The planar interference graph used in the proof of Theorem 9.

It can be easily seen that the expected benefit of the optimal off–line algorithm on sequences of calls generated according to $\mathcal{P}$ is

$$\mathcal{E}_{\mathcal{P}}[B_{OPT}(\sigma)] = 1 \cdot \frac{4}{5} + 5 \cdot \frac{1}{5} \cdot \frac{2}{7} + 7 \cdot \frac{1}{5} \cdot \frac{5}{7} = \frac{73}{35}.$$

Let A be a deterministic call control algorithm that runs on the calls produced by $\mathcal{ADV}$. 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 c_0, and q_1 the number of executions in which A accepts both calls in cells c_1 and c_2.

The expected number of executions in which the algorithm does not accept the call in cell c_0 and the adversary produces calls in cells $c_1, ..., c_5$ is $\frac{1}{5}(t - q_0)$. Similarly, the expected number of executions in which the algorithm does not accept the calls in cells $c_0, c_1, ..., c_5$ and the adversary produces calls in cells $c_6, ..., c_{12}$ is $\frac{5}{7}\left(\frac{1}{5}(t - q_0) - q_1\right)$. Thus,

$$\mathcal{E}_{\mathcal{P}}[B_A(\sigma)] \leq \frac{q_0 + 5q_1 + 7\frac{5}{7}\left(\frac{1}{5}(t - q_0) - q_1\right)}{t} = 1,$$

and

$$c_A^{\mathcal{P}} \geq \frac{73}{35} = 2.086.$$

By Lemma 7, we obtain that this is a lower bound on the competitive ratio of any randomized algorithm against an oblivious adversary. □

4. 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] 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.

[3] I. Caragiannis, C. Kaklamanis, and E. Papaioannou. On–line Call Control in Cellular Networks. *Foundations of Mobile Computing* (satellite workshop of FST&TCS '99), 1999. Also, CTI Technical Report 99-10-02.

[4] W.K. Hale. Frequency Assignment: Theory and Applications. In *Proceedings of the IEEE*, 68(12), pp. 1497–1514, 1980.

[5] 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.

[6] 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.

[7] C. McDiarmid and B. Reed. Channel Assignment and Weighted Coloring. Manuscript, 1997.

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

[9] L. Narayanan and S. Shende. Static Frequency Assignment in Cellular Networks. In *Proceedings of the 5th International Colloquium on Structural Information and Communication Complexity (SIROCCO '97)*, pp. 215–227, 1997.

[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] 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.