DIMACS Series in Discrete Mathematics
and Theoretical Computer Science

Frequency Assignment in Mobile and Radio Networks

Dimitris Fotakis, Grammati Pantziou, George Pentaris, and Paul Spirakis

Abstract. We deal with the problem of frequency assignment in mobile and general radio networks, where the signal interferences are modeled using an interference graph G. Our approach uses graph theoretic and optimization techniques. We first study on-line algorithms for frequency assignment in mobile networks. We prove that the greedy algorithm is Δ-competitive, where Δ is the maximum degree of G. We next employ the "classify and randomly select" paradigm to give a 5-competitive randomized algorithm for the case of planar interference graphs. We also show how the problem of on-line frequency assignment in mobile networks with multiple available frequency channels reduces to the problem of on-line frequency assignment in mobile networks with a single channel.

We continue to study *radio coloring* and *radio labeling* as combinatorial models for frequency assignment in general radio networks. In both problems, the objective is to minimize the maximum frequency channel used, while the transmitters being adjacent in the interference graph must be assigned channels that differ by at least two from each other. In radio coloring, different channels must be assigned to transmitters that are at distance two in the interference graph. Additionally, in radio labeling, all the transmitters must be assigned distinct frequency channels. Radio labelling is shown to be equivalent to a generalization of Hamiltonian path, and both problems remain $\mathcal{NP}$-complete, even if they restricted to graphs of diameter two. We finally present algorithms and lower bounds for two on-line variations of radio labeling.

1. Introduction

The problem of *frequency assignment* arises when different radio *transmitters* which operate in the same geographical area interfere to each other when assigned to the same or closely related frequency channels. This situation is common in a wide variety of real world applications related to *mobile* or general *radio networks*. The most common model for an instance of the Frequency Assignment Problem (FAP) is the *interference graph*. Each vertex of an interference graph represents a transmitter, while each edge represents an interference constraint between two adjacent transmitters. Additionally, the available frequency channels are assumed to be uniformly spaced in the spectrum, and are labelled using positive integers.

1991 *Mathematics Subject Classification.* Primary 68Q25, 68R10; Secondary 05C15, 90D43.

All the authors were partially supported by ESPRIT LTR Project no. 20244—ALCOM-IT.

Frequency channels with adjacent integer labels are also assumed adjacent in the spectrum [**GL, WR**].

In mobile (cellular) networks, the moving hosts usually communicate by modulation of radio waves. The geographical area that is served by a cellular communication network is divided into smaller regions called *cells*. Each cell has a base station, also referred to as the *mobile service station*. Each mobile service station is connected to the service stations of the neighboring cells by a fixed high bandwidth network. To establish a communication session, or place a call, a mobile user, also referred to as *mobile host*, has to send a request to the service station of the cell to which it currently belongs. The call can be supported only if one (possibly more, depending on the required bandwidth) wireless channel can be allocated for communication between the mobile host and the mobile service station.

Wireless channels form divisions of the frequency spectrum; the signal carried on a wireless channel is carried on the respective frequency. A particular wireless channel cannot be used to support more than one call in a cell or neighboring cells since the calls will interfere to each other; this is called *channel or signal interference*. On the other hand, the same wireless channel can be used to support calls in geographically distant cells. An efficient channel (or frequency) allocation strategy in a mobile network should guarantee that no signal interference is taking place, and exploit the principle of frequency reuse to increase the availability of wireless channels.

Under these considerations, the problem of channel assignment in mobile networks can be viewed as a *graph multicoloring problem* [**KK, NS**], where each vertex v of the interference graph has an associated integer weight $w(v)$ representing the number of calls that must be served by the corresponding base station. We seek to assign $w(v)$ distinct colors to each vertex of the interference graph, such that for each vertex $\{v, u\}$ the set of colors assigned to the endpoints v and u are disjoint. The objective is to minimize the number of different colors used by such an assignment.

In a general radio network, the geographical regions covered by the radio transmitters are usually intersected with each other. Suppose that a radio receiver is "tuned" to a signal on a frequency channel c_0, broadcast by its local transmitter, that is the one at closest distance. Reception will be degraded if excessive interference is created by other transmitters in the vicinity. First, there is "co-channel" interference, due to reuse of the frequency channel c_0 at nearby sites; there are also contributions from sites using channels close to c_0, since, in practice, neither transmitters nor receivers operate exclusively within the frequency channels that they are assigned [**GL, HLS**].

If we only consider "co-channel" interference constraints between adjacent transmitters, the problem of frequency assignment in radio networks is equivalent to the problem of coloring the interference graph with the minimum number of colors. In order to ensure acceptable signal quality, constraints need to be imposed on the allowed channel separations between pairs of potentially interfering transmitters [**GL, Hale, Rayc**]. Moreover, if the transmitters are assumed to cover a local or metropolitan area, the frequency channels are non-reusable in the sense that each channel must be assigned to at most one transmitter.

In relation to radio networks, we study two combinatorial problems modelling "co-channel" and "adjacent-channel" interference constraints: *radio coloring* and *radio labeling*. In both of them, the objective is to minimize the maximum channel used, while adjacent transmitters must be assigned non-adjacent channels, i.e. channels differing by at least two from each other. In radio coloring, we also impose the constraint that transmitters at distance two are not allowed to use the same channel. However, the same channel can be used by a set of transmitters, that are at distance more than two from each other, while in radio labeling all transmitters must be assigned distinct channels/labels.

1.1. Summary of Results. We first deal with the competitive analysis of *on-line algorithms* for frequency assignment in mobile networks. We study general interference graphs and we prove that the greedy algorithm is Δ-competitive, where Δ is the maximum degree of the interference graph. We next prove a lower bound of Δ on the competitive ratio of on-line algorithms which do not reject requests by choice. Then, we study the case of *planar* interference graphs and we employ the "classify and randomly select" paradigm to give a 5-competitive randomized algorithm. This paradigm has been used in previous work concerning competitive analysis in order to construct randomized competitive algorithms for benefit problems [**ABFR, LT**]. We finally show how the problem of on-line frequency assignment in mobile networks with multiple available frequency channels reduces to the problem of on-line frequency assignment in mobile networks with a single channel.

We proceed to deal with general radio networks by proving that radio labelling is equivalent to Hamiltonian Path with distances one and two (HP(1,2)). Therefore, the results of [**PY3**] imply that radio labeling is MAX–SNP-hard and approximable in polynomial time within 7/6. Additionally, we show that both radio coloring and radio labelling remain $\mathcal{NP}$-complete, even if they are restricted to graphs of diameter two. We finally define two on-line variations of radio labeling, and we prove that the greedy algorithm achieves competitive ratio 2 for both variations. We also provide a lower bound of 3/2 on the performance of any on-line algorithm for the first variation against any adversary.

1.2. Outline of the Paper. In Section 2, we give some preliminary definitions and technical lemmas, as well as the definitions of the problems we deal with. We proceed to an overview of frequency assignment techniques for mobile networks in Section 3, and in Section 3.1 we present the previous work concerning this topic. In Section 3.2, we deal with general interference graphs proving that the greedy algorithm is Δ-competitive. We next obtain a lower bound for algorithms which do not reject by choice. In Section 3.3 we deal with planar interference graphs and in section 3.4 we present the reduction from many frequency channels to a single channel. In Section 4.1, we prove that radio labelling is equivalent to HP(1,2) and both radio coloring and radio labelling remain $\mathcal{NP}$-complete for graphs of diameter two. In Section 4.2, we analyze the performance of greedy algorithms for on-line radio labelling, and we obtain a lower bound on the competitive ratio of any on-line algorithm. Some concluding remarks and directions for further research are discussed in Section 5.

2. Definitions and Techniques

Given a graph $G(V, E)$, let Δ denote the maximum degree of G, $d(v, u)$ denote the length of the shortest path between $v, u \in V$, and $\text{diam}(G)$ be the diameter of the graph G, defined as $\text{diam}(G) = \max_{v,u \in V}\{d(v,u)\}$. Also, the complementary graph of G, denoted $\overline{G}$, is a graph on the same vertex set V that contains an edge $\{v,u\}$ if and only if $\{v,u\} \notin E$.

Hamiltonian Path with distances one and two (HP(1,2)) is the problem of finding a Hamiltonian path of minimum length in a complete graph, where all the edge lengths are either one or two. A Hamiltonian path is a simple path visiting each vertex of a graph exactly once. An instance of HP(1,2) can also be defined implicitly by an unweighted graph $G(V,E)$, if the edges of E are considered of length one, and the edges not in E (non-edges) of length two. Therefore, HP(1,2) is a generalization of the Hamiltonian path problem. HP(1,2) is MAX–SNP-hard [**PY1**] and approximable in polynomial time within a factor of 7/6 [**PY3**].

2.1. Competitive Analysis and Lower Bounds. The usual approach of analyzing the performance of on-line algorithms is the *competitive analysis* [**ST**]. In competitive analysis, the performance of the on-line algorithm is compared to the performance of the optimal off-line algorithm on every sequence of requests, and the worst-case ratio is considered. Let A be an on-line algorithm for a benefit (maximization) problem and σ any sequence of requests. Also, let $A(\sigma)$ denote the benefit accrued by A when presented with σ, and $OPT(\sigma)$ denote the benefit accrued by the optimal off-line algorithm OPT on the same sequence. We say that the on-line algorithm A is *c-competitive* if there exists a constant b such that on every request sequence σ,

$$c \cdot A(\sigma) + b \geq OPT(\sigma)$$

The competitive ratio for cost (minimization) problems is defined similarly. In case of randomized on-line algorithms, the competitive ratio is defined with respect to the expected benefit of the algorithm.

A standard technique used in competitive analysis is to employ an *adversary* which plays against the algorithm A and constructs an input which incurs a high cost for A and a low cost for OPT. We say that an adversary is *oblivious* if it constructs the request sequence in advance, before A starts working on the sequence. However, the oblivious adversary knows the probability distribution of the actions taken by the algorithm A. An adversary is *non-oblivious* if it decides the next request based on the on-line algorithm's answers to previous requests. Notice that in case of deterministic algorithms, the non-oblivious adversary is equivalent to the oblivious one since the on-line algorithm's answers are completely predictable.

A common technique for proving lower bounds on the competitive ratio of a randomized on-line algorithm against the oblivious adversary is to bound the performance of the "best" deterministic on-line algorithm on inputs generated from the "worst" probability distribution. In particular, let $\mathcal{P}$ be a probability distribution over sequences of requests σ, let $\mathrm{E}_{\mathcal{P}}[A(\sigma)]$ denote the expected value of the benefit of an algorithm A, and $\mathrm{E}_{\mathcal{P}}[OPT(\sigma)]$ be the expected benefit of the optimal off-line algorithm on inputs generated from $\mathcal{P}$. An algorithm A (for a benefit problem) is c-competitive against $\mathcal{P}$ if there exists a constant b such that,

$$c \cdot \mathrm{E}_{\mathcal{P}}[A(\sigma)] + b \geq \mathrm{E}_{\mathcal{P}}[OPT(\sigma)]$$

The following theorem can be derived by Yao's Lemma based on the minimax principle of game theory [**Yao**].

LEMMA 2.1 ([**BLS**]). *A real number c is a lower bound on the competitive ratio of randomized on-line algorithms against the oblivious adversary if and only if there exists a probability distribution $\mathcal{P}$ such that c is a lower bound on the competitive ratio of any deterministic on-line algorithm against $\mathcal{P}$.*

2.2. Mobile Networks. The interference among signals can be modeled as an *interference graph*. The base stations (transmitters) are represented by the vertices of the interference graph, while the edges denote spatial adjacency and potential signal interference between the connected base stations. In particular, two base stations connected by an edge are not allowed to operate at the same frequency channel in order to avoid co-channel interference. In the following, we always denote the interference graph by $G(V, E)$, $|V| = n$.

In case of mobile networks, each vertex v_i, $i \in \{1, \cdots, n\}$, is associated with a set of colors C_i, representing the set of free frequency channels from which the channels to be allocated to mobile hosts by v_i can be chosen. A request r_i is a tuple $< v_j, t_i, d_i, b_i >$, where $v_j \in V$ represents the base station the request is applied to, t_i denotes the arrival time of the request, d_i denotes the duration or the service time of the request, and b_i is the benefit accrued by the algorithm if it accepts the request. Requests are presented to a call admission control algorithm in an on-line manner and the algorithm must, upon receipt, either accept or reject the request.

A vertex v_j can accept a request r_i applied to it by assigning a channel/color from the set C_j to r_i, in case there is an available color in C_j. Once a request is accepted, it continues to use the frequency channel allocated to it for d_i time units until completion; that is, no preemption is allowed. During that time period, the color assigned to r_i cannot be used to serve any other request applied to either v_j or the neighbors of v_j in G. We focus on the most basic case where the durations of the requests are infinite, and all benefits associated with the requests are the same; without loss of generality we assume that these benefits have unit value.

2.3. General Radio Networks. In order to avoid signal interference in general radio networks, the spatially adjacent transmitters should operate not only at different but also at separated frequency channels. The interference graph usually represents the constraints providing a minimum allowed separation between the channels assigned to each pair of transmitters. Hence, graph coloring provides a very limited model for frequency assignment in radio networks [**AHLT, Hale, Rayc**].

A *valid coloring* is an assignment of *colors* to the vertices of a graph such that no two adjacent vertices get the same color. In the graph coloring problem, the objective is to minimize the number of different colors used by a valid coloring. The *chromatic number* $\chi(G)$ of a graph G is the number of colors used by an *optimal* valid coloring. The decision version of graph coloring is $\mathcal{N}P$-complete [**GJ**], while it is $\mathcal{N}P$-hard to efficiently approximate the chromatic number [**LY**]. However, there exist polynomial-time approximation algorithms whose performance guarantees are not so good [**Hall, KMS**].

2.3.1. *Graph Radio Coloring.* Instead of specifying channel separation constraints for each pair of transmitters, the frequency assignment problem can be defined by specifying a minimum allowed distance for each spectral separation that is a potential source of interference. In this way, the interference constraints can be given as a set of distances $\{D_0, D_1, \ldots, D_\kappa\}$, where the distance D_x, $0 \le x \le \kappa$, gives the minimum distance for the use of channels/labels at distance x (see also

[**GL, Hale**]). The following generalized version of graph coloring is derived by applying the aforementioned consideration to unweighted graphs.

DEFINITION 2.2. (**κ-coloring**)
INPUT: A graph $G(V, E)$.
GOAL: Compute a valid κ-coloring of minimum value ν, that is a function $\chi_\kappa : V \mapsto \{1, \dots, \nu\}$ such that $|\chi_\kappa(v) - \chi_\kappa(u)| = x$, for some $x = 0, 1, \dots, \kappa$, only if $d(v, u) \geq k - x + 1$.

Clearly, the κ-coloring problem allows a pair of non-adjacent transmitters to be assigned the same channel, provided they are located far apart. In the sequel, we shall concentrate on the study of 2-coloring, also called *radio coloring*, modelling the "co-channel" and "adjacent-channel" interference constraints. The definition of radio coloring is implicit in many previous works [**Hale, Rayc**].

We remark that the problem of radio coloring the graph G is not equivalent to the problem of coloring the square graph[1] of G, denoted G^2. In particular, in a graph coloring of G^2, the objective is to minimize the number of different colors used, while in a radio coloring of G, the objective is to minimize the largest color/label assigned to a vertex, called color span. For example, if G is $K_{m,m}$, which is the complete bipartite graph with m vertices on each class, then G^2 is the complete graph on $2m$ vertices. Therefore, $\chi(G^2) = 2m$, while the radio chromatic number of G is $2m + 1$.

2.3.2. *Graph Radio Labelling.* The definition of κ-labeling was communicated to us by [**Hara**]. The κ-labelling problem is also defined by a set of distances $\{D_1, \dots, D_\kappa\}$, where the distance D_x, $1 \geq x \geq \kappa$, gives the minimum distance for the use of channels/labels at distance x. The κ-labelling problem can be thought as an instance of the frequency assignment problem with the minimum distance for channel re-use being infinite, i.e. $D_0 = \infty$. Hence, channel re-use is not allowed, and all the transmitters must be assigned distinct labels.

DEFINITION 2.3. (**κ-labeling**)
INPUT: A graph $G(V, E)$.
GOAL: Compute a valid κ-labelling of minimum value ν, that is a function $\lambda_\kappa : V \mapsto \{1, \dots, \nu\}$ such that $|\lambda_\kappa(v) - \lambda_\kappa(u)| = x$, for some $x = 1, \dots, \kappa$, only if $d(v, u) \geq k - x + 1$.

In the sequel, we shall concentrate on the 2-labelling problem, also called *radio labeling*, which imposes "co-channel" interference constraints among all transmitter pairs, and "adjacent-channel" constraints among pairs of adjacent transmitters.

Given a graph $G(V, E)$, let $\lambda_2(G)$ denote the value of an optimal radio labelling for G. Obviously, given a coloring of G with χ colors, one can easily compute a radio labelling of value $|V| + \chi - 1$, since an assignment of contiguous labels to the vertices of the same color class is a (partial) valid radio labelling. Therefore,

$$|V| \leq \lambda_2(G) \leq |V| + \chi(G) - 1$$

It is not hard to verify that $\lambda_2(G) = |V|$ for all graphs G such that $\overline{G}$ contains a Hamiltonian path. On the other hand, $\lambda_2(G) = |V| + \chi(G) - 1$ for any complete r-partite graph G, $r \geq 2$.

[1]Two vertices are connected in G^2, if they are at distance at most two in G.

2.3.3. *On-line Radio Labelling.* Motivated by the problem of on-line frequency assignment in mobile and general radio networks, we proceed to define the problem of on-line radio labeling. In the on-line setting of the problem, an induced subgraph of the interference graph appears to the on-line algorithm in a vertex by vertex fashion. In case of on-line radio labelling, a request is a newly arrived vertex that has to be assigned a radio label. A request is accepted if the vertex actually gets a valid label. Otherwise, the request is rejected. We define two variations of on-line radio labelling. The first variation is motivated by call admission control, while the second one is motivated by on-line optimization.

DEFINITION 2.4. **(On-line Radio Labelling with Bound)**
INPUT: A graph $G(V, E)$ and an integer bound $\beta > 0$ are known to the algorithm. A new vertex $v \in V$ is presented to the on-line algorithm at every step. The adversary may choose to request any subset $V' \subseteq V$ in any order. The algorithm has to decide to either accept v by assigning to it a label no greater than β or reject it. At any step, the labels of the set V_a of accepted vertices must form a valid radio labelling for the subgraph of G induced by V_a.
GOAL: Maximize the number of accepted vertices.

DEFINITION 2.5. **(Radio Labeling without Bound)**
INPUT: A graph $G(V, E)$ is known to the algorithm. At every step, a new vertex $v \in V$ is presented to the on-line algorithm, and it is immediately assigned a label. The adversary may choose to request any subset $V' \subseteq V$ in any order. At any step, the labels of the set V_p of the vertices presented so far must form a valid radio labelling for the subgraph of G induced by V_p.
GOAL: Minimize the maximum label used.

3. Frequency Assignment in Mobile Networks

Channel allocation algorithms can be classified as *fixed channel*, where the set of channels allocated to a cell cannot change with time, or *dynamic channel* where the set of channels allocated to a cell may vary with time. The dynamic channel allocation approach is more flexible, since it allows easier adaptation to the number of requests and bursts of the network usage that may appear in particular areas.

Another classification of the channel allocation algorithms is based on the locality of the proposed solution. Some of them are *centralized*, which means that one base station collects the requests from all the mobile service stations and decides how to satisfy them; in *distributed* solutions, however, the allocation problem is solved locally by the mobile service stations of the involved cells and their neighbors with the use of a distributed protocol. The distributed approach is more robust and can tolerate system failures more successfully than the centralized one.

3.1. Previous work. In the context of a cell and its neighbors, the use of a particular channel to support a communication session is equivalent to an execution of a critical section by the cell in which the channel is being used. Several neighboring cells may be concurrently trying to choose channels to support sessions in their region; this may lead to conflicts due to the limited number of available channels. The resolution of such conflicts resembles the mutual exclusion problem in distributed systems.

However, the channel allocation problem is different than the mutual exclusion problem. First, a cell may be supporting multiple communication sessions from different mobile hosts in its region using a different channel for each session. This is equivalent to a cell being in multiple distinct critical sessions concurrently. Additionally, existing mutual exclusion algorithms for distributed systems [**CM, CS, GK**] assume that a node specifies the identity of the resource it wants to access in a critical section; in contrast, in the channel allocation problem a cell asks for *any* available channel.

A distributed approach to channel allocation problem is presented in [**PPS**]. When a mobile host wants to establish a new communication session, it initiates a channel allocation protocol by presenting a request to its mobile service station. The mobile service station executes the protocol by exchanging messages with its neighboring mobile service stations to determine if the request can be satisfied. If so, it allocates a communication channel for the mobile host to use for the duration of the session. The protocol requires for each request $O(M)$ messages to be exchanged between the mobile service station and its neighbors, where M is the number of those neighbors.

The problem of allocating channels to mobile hosts by the base stations can be also analyzed as a generalized version of the *list coloring* problem [**AT, ERT, JT**]. In the list coloring problem, every vertex of the graph has associated with it a list of colors which, in our case, represent the set of free channels. The requirement is to find a proper vertex-coloring of the graph such that each vertex is colored with one of the colors in its list. Thus, any protocol for the list coloring problem that can be invoked multiple times with lists and requests varying in time, can be used for the channel allocation problem in a dynamic manner.

In addition to the performance measures described above, another important property of the solution is the ability to cope with failures, since such are unavoidable in any distributed setting. One criterion to measure this property is the *failure locality*, which measures the size of the network affected by the faulty station. The distributed list coloring protocol presented in [**GPT**] achieves failure locality 4, and requires for each request $O(\Delta^2)$ messages to be exchanged between the mobile service stations for general interference graphs, and $O(\Delta)$ messages for planar interference graphs.

3.2. General graphs. We assume that all colors are initially available at all vertices of the interference graph. That is, $C_1 = C_2 = \cdots = C_n$, $|C_1| = |C_2| = \cdots = |C_n| = c$. In this section, we show that a competitive ratio of Δ is achievable by a simple algorithm. The algorithm does not *reject by choice*; that is, it always accepts a request applied to a vertex v_j if there is an available color in C_i, and therefore a free frequency channel, to serve the request. We also show that any algorithm that decides not to reject by choice, has a competitive ratio no less than Δ.

Consider the following greedy on-line algorithm A: Given a request sequence $\sigma = r_1, r_2, \ldots, r_N$, A accepts r_i applied to v_j at time t_i, by assigning to r_i a color from the set C_j that is available at time t_i. If such a color does not exist, A rejects r_i. We remark that the distributed protocol proposed in [**PPS**] is the implementation of the algorithm A in a distributed environment.

THEOREM 3.1. *Algorithm A is Δ-competitive.*

PROOF. Let AC_{OPT} be the set of requests accepted by the optimal off-line algorithm. Let RJ_{OPT} be the set of requests rejected by the optimal off-line algorithm because of requests in AC_{OPT}. Now, let AC_A and RJ_A be respectively, the corresponding sets of accepted and rejected requests by the greedy algorithm A. Let $ac_{OPT} = |AC_{OPT}|$, $rj_{OPT} = |RJ_{OPT}|$, $ac_A = |AC_A|$, and $rj_A = |RJ_A|$. Note that $N = ac_{OPT} + rj_{OPT}$ (N is the size of the input request sequence σ), and $N = ac_A + rj_A$. Suppose that $ac_A = k_1 + k_2$ where k_1 is the cardinality of the set

$$\{r \mid r \in AC_A \ \ and \ \ r \in AC_{OPT}\}$$

including requests accepted by both the optimal and A ("optimal" requests), and k_2 is the cardinality of the set

$$\{r \mid r \in AC_A \ \ and \ \ r \in RJ_{OPT}\}$$

which contains requests that have been rejected by the optimal and are accepted by A ("non-optimal" requests). The accepted k_1 "optimal" requests may cause the rejection of at most $rj_{OPT} - k_2$ other requests. This happens because any "optimal" request can cause the rejection of all the requests that are in RJ_{OPT} except the accepted k_2. Note that any "optimal" request cannot cause the rejection of any other request that is in AC_{OPT}. Furthermore, A can reject at most $k_2\Delta$ requests because of the k_2 "non-optimal" requests that accepts. This holds because any "non-optimal" request may cause the rejection of at most Δ "optimal" requests (requests that are in AC_{OPT}). Since the total number of the requests rejected by A is the sum of requests rejected because of the accepted k_1 and k_2 we have that

$$rj_A \leq rj_{OPT} - k_2 + k_2\Delta$$

Since $ac_A = N - rj_A$ it follows that

$$\begin{aligned}
ac_A &\geq N - rj_{OPT} - k_2(\Delta - 1). \text{ But } \ N - rj_{OPT} = ac_{OPT}, \ \text{ so} \\
ac_A &\geq ac_{OPT} - k_2(\Delta - 1). \text{ Since } \ k_2 = ac_A - k_1, \ \text{ we have that} \\
ac_A &\geq ac_{OPT} - (ac_A - k_1)(\Delta - 1). \text{ Then} \\
ac_A\Delta &\geq ac_{OPT}, \ \text{ and finally} \\
\frac{ac_{OPT}}{ac_A} &\leq \Delta
\end{aligned}$$

Therefore, A is Δ-competitive. □

3.2.1. *Lower Bound.* We show a corresponding lower bound for on-line algorithms that do not decide to reject by choice.

LEMMA 3.2. *Any on-line algorithm that decides not to reject by choice, has a competitive ratio at least Δ.*

PROOF. Let A' be an algorithm that does not reject by choice. We can always construct a request sequence σ that leads the algorithm to a competitive ratio equal to Δ.

Let v_j be a vertex of maximum degree, and for each pair w_i, w_j of its neighbors there is no edge between w_i and w_j. Consider a request sequence σ which starts at time 0 with a set of c requests applied to vertex v_j, and continues at time 1 with c requests to each one of the neighbors of v_j. Note that A' serves only c requests, while an optimal off-line algorithm may serve Δc. □

3.3. Planar Graphs. In this section we apply the *"classify and randomly select"* paradigm [**AAFLR, ABFR, LT**] and give a randomized 4-competitive channel allocation algorithm for the case that the interference graph is planar. According to this paradigm, the randomized algorithm classifies the requests applied to the graph into a number of classes. It then randomly selects one of the classes, and considers only requests that belong to the selected class, rejecting all other requests. In our case the classification of the requests is done with respect to which of the vertices of G each request is applied to. Toward this end, a preprocessing of the graph is done, which produces a partition of its vertices into a number of classes.

Specifically for the case of planar graphs, our classification algorithm is a 5-coloring algorithm that partitions the planar graph into 5 classes. The randomized on-line algorithm B selects uniformly at random one of the five classes and considers only requests at vertices that belong to this class. Once such a request at a vertex v is received, the greedy algorithm is used, and the request is accepted if no other request at v has been accepted before. We show:

LEMMA 3.3. *Algorithm B for planar graphs is 5-competitive against an oblivious adversary.*

PROOF. Let OPT_i and A_i be the benefit of the optimal off-line algorithm and the greedy algorithm, respectively, restricted to requests for the vertices of class i, $i \in \{1, 2, 3, 4, 5\}$. Since there is no edge connecting two vertices at the same class, we have $OPT_i = A_i$. Let now OPT be the number of requests accepted by the optimal off-line solution in the whole network. Since each request accepted by the optimal algorithm belongs to some class,

$$OPT \leq \sum_{i=1}^{5} OPT_i$$

The on-line algorithm selects uniformly at random one among the 5 classes and obtains the optimal benefit for that class. Therefore, the expected benefit of the on-line algorithm is

$$\sum_{i=1}^{5} \frac{1}{5} OPT_i \geq \frac{1}{5} OPT$$

It follows that the on-line algorithm is 5-competitive. □

Since a tree and a graph with maximum degree Δ allow for a 2-coloring and a $(\Delta + 1)$-coloring respectively, it immediately follows:

COROLLARY 3.4. *There is an on-line randomized algorithm for trees which is 2-competitive against an oblivious adversary.*

COROLLARY 3.5. *There is an on-line randomized algorithm for graphs of maximum degree Δ which is $(\Delta + 1)$-competitive against an oblivious adversary.*

3.4. Reduction from many frequency channels to one. We show that the problem of on-line frequency assignment in mobile networks with multiple available frequency channels reduces to the problem of on-line frequency assignment in networks with a single channel. The technique used is due to Awerbuch *et al.* [**AAFLR**], and it can be applied to any on-line benefit problem, where the benefit is gained by accommodating entities in any of several independent *bins*.

LEMMA 3.6 ([**AAFLR**]). *Let A_c be a c-competitive call control algorithm for frequency allocation in mobile networks with one frequency channel. Then, there exists a $(c+1)$-competitive call control algorithm for frequency allocation in mobile networks with any number of frequency channels.*

Using this reduction, an on-line randomized algorithm concerning general networks with a single frequency channel is obtained in [**PPS**]. This algorithm is $O(\delta)$-competitive where δ is the average degree of G.

4. Frequency Assignment in Radio Networks

4.1. The Complexity of Radio Coloring and Radio Labelling. Since radio labelling assigns distinct integer labels to all the vertices of the interference graph, it is a vertex arrangement problem. In particular, we show that radio labeling is equivalent to HP(1,2) in the complementary graph.

THEOREM 4.1. *Graph radio labeling and HP(1,2) are equivalent.*

PROOF. Given an instance of radio labeling, i.e. a graph $G(V,E)$, the corresponding instance of HP(1,2) is a complete graph $\hat{G}$ on the vertex set V, and the distance function $\hat{d}$ is defined for all $v,u \in V$, $v \neq u$, by

$$\hat{d}(v,u) = \begin{cases} 1 & \text{if } (v,u) \notin E \\ 2 & \text{if } (v,u) \in E \end{cases}$$

Given any valid radio labelling L (for G) of value $\lambda_2(L)$, we can obtain a Hamiltonian path H for $\hat{G}$ by traversing all the vertices in increasing order of their labels. Moreover, the following claim implies that the length of the Hamiltonian path H is $\lambda_2(L)-1$.

CLAIM 1. The length of the path up to any vertex of label i, $i = 1, \ldots, \lambda_2(L)$, is exactly $i-1$.

PROOF. We prove the claim by induction on i. Clearly, it is true for the first vertex, where $i=1$. Assume inductively that it is true for any vertex v of label $i \geq 1$, and let u be the next vertex in the Hamiltonian path. We proceed by case analysis:

1. If the label of u is $i+1$, the edge $\{v,u\}$ is not present in G and, by construction, $\hat{d}(v,u) = 1$. Thus, the length of the path up to vertex u is exactly i.
2. If the label of u is $i+2$, by the construction of the Hamiltonian path, there does not exist a vertex of label $i+1$. Therefore, the edge $\{v,u\} \in E$, and $\hat{d}(v,u) = 2$. Consequently, the path up to u has length exactly $i+1$.

□

Conversely, given an instance $(\hat{G}, \hat{d})$ of HP(1,2) on the vertex set V, $|V| = n$, an instance $G(V,E)$ of radio labeling can be obtained by only connecting the vertex pairs that are at distance 2. Furthermore, given a Hamiltonian path $H = (v_1, v_2, \ldots, v_n)$ for $\hat{G}$ of length $l(H)$, we obtain a valid radio labelling L (for G) as follows:

1. $\lambda_2(v_1) = 1$.
2. For $i = 1, \ldots, n-1$,

(a) $\lambda_2(v_{i+1}) = \lambda_2(v_i) + 1$, if $\hat{d}(v_i, v_{i+1}) = 1$. By the construction of $G(V, E)$, in this case $\{v_i, v_{i+1}\} \notin E$.
(b) $\lambda_2(v_{i+1}) = \lambda_2(v_i) + 2$, if $\hat{d}(v_i, v_{i+1}) = 2$. By the construction of $G(V, E)$, in this case $\{v_i, v_{i+1}\} \in E$.

By construction, all the vertices get distinct labels, while, if an edge $\{v_i, v_{i+1}\}$ is present in E, the vertices v_i and v_{i+1} are assigned non-adjacent labels. Therefore, the resulting radio labeling is a valid one. Additionally, the last vertex of the path v_n is assigned the label $l(H)+1$, that is the largest label used. Hence, $\lambda_2(L) = l(H)+1$. □

The previous theorem implies that radio labeling is $\mathcal{N}P$-complete, since it is equivalent to a generalization of the Hamiltonian Path problem. Additionally, HP(1,2) was shown MAX–SNP-hard and approximable in polynomial time within $\frac{7}{6}$ in [**PY3**]. Therefore, Theorem 4.1 implies that radio labelling is also MAX–SNP-hard and approximable within $\frac{7}{6}$. Furthermore, the following is a straight forward consequence of Thereom 4.1.

COROLLARY 4.2. *For any graph $G(V, E)$, $\lambda_2(G) \leq |V|$ if and only if the complementary graph $\overline{G}$ contains a Hamiltonian path.*

PROOF. The proof of Theorem 4.1 implies that, for any graph $G(V, E)$, there exists a radio labelling of value at most $|V|$ if and only if the corresponding instance of HP(1,2) contains a Hamiltonian path of length $|V| - 1$, that is a Hamiltonian path entirely consisting of edges of length 1. By the construction of the HP(1,2) instance, an edge e has length 1 if and only if $e \notin E$, i.e. if and only if e is present in the complementary graph $\overline{G}$. □

We proceed to show that both radio labelling and radio coloring remain $\mathcal{N}P$-complete, even for a very restricted class of instances, namely graphs of diameter two.

THEOREM 4.3. *Radio labelling and radio coloring restricted to graphs of diameter two are $\mathcal{N}P$-complete.*

PROOF. Let $G(V, E)$, $|V| = n$, be any graph of diameter two. Since, for all $v, u \in V$, $d(v, u)$ is at most two, any valid radio coloring must assign distinct integer colors to all the vertices of G. Moreover, if $\{v, u\} \in E$, then $|\chi_2(v) - \chi_2(u)| \geq 2$. Therefore, if diam$(G) = 2$, the problem of radio coloring the graph G is equivalent to the problem of radio labelling G. Hence, it suffices to show that radio labelling is $\mathcal{N}P$-complete for graphs of diameter two.

Clearly, radio labelling is in $\mathcal{N}P$. Additionally, Corollary 4.2 implies that $\lambda_2(G) \leq |V|$ if and only if the complementary graph $\overline{G}$ contains a Hamiltonian path. Thus, in order to show that radio labelling is $\mathcal{N}P$-complete for graphs G with diameter two, it suffices to show that the Hamiltonian Path problem remains $\mathcal{N}P$-complete for graphs $\overline{G}$, that are complements of graphs having diameter two.

Let $G'(V', E')$ be any graph and let $s, t \in V'$ be any non-adjacent vertex pair. The problem of deciding if G' contains a Hamiltonian path starting from s and ending to t is $\mathcal{N}P$-complete (HAMILTONIAN PATH BETWEEN TWO VERTICES [**GJ**]). Let $G^{(c)}(V' \cup \{v_s, v_t\}, E' \cup \{(s, v_s), (t, v_t)\})$ be the graph obtained from G' by adding two non-adjacent vertices v_s, v_t, and connecting v_s to s and v_t to t (Figure 1). The graph $G^{(c)}$ is the complement of a graph of diameter two. In

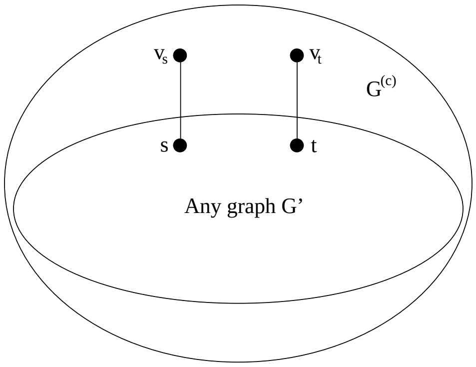

FIGURE 1. The complementary graph of $G^{(c)}$ has diameter two.

particular, the following observations justify that $\mathrm{diam}(\overline{G^{(c)}}) = 2$, since all the vertices are at distance at most two from each other in $\overline{G^{(c)}}$.

1. The vertex pairs $(s,t), (v_s,v_t), (s,v_t)$ and (t,v_s) are connected by edges.
2. Any pair of vertices $u, w \in V' - \{s,t\}$ are at distance at most two from each other, because they are connected to both v_s, v_t.
3. Any vertex $u \in V' \cup \{v_s\} - \{s,t\}$ is at distance at most two from s, because both u and s are connected to v_t.
4. Any vertex $u \in V' \cup \{v_t\} - \{s,t\}$ is at distance at most two from t, because both u and t are connected to v_s.

Additionally, $G^{(c)}$ contains a Hamiltonian path if and only if G' contains a Hamiltonian path from s to t. Therefore, Hamiltonian path is $\mathcal{NP}$-complete for complements of graphs of diameter two. □

4.2. Algorithms and Lower Bounds for On-line Radio Labeling.

4.2.1. *Radio Labeling with Bound.* We first analyze the performance of the greedy algorithm BGREEDY. BGREEDY assigns to a newly arrived request the least integer j, such that both j has not been assigned to any previously accepted request, and the resultant radio labeling is a valid one. BGREEDY accepts the request, if and only if j is no more than the bound β.

LEMMA 4.4. *The competitive ratio of* BGREEDY *is* $\beta \left\lceil \frac{\beta}{2} \right\rceil^{-1}$.

PROOF. The BGREEDY algorithm always accepts at least $\left\lceil \frac{\beta}{2} \right\rceil$ requests, because, for $i = 1, 2, \ldots, \left\lceil \frac{\beta}{2} \right\rceil$, BGREEDY can always assign a label no more than $2i-1$ to the i-th request. Since the optimal off-line algorithm cannot accept more than β requests, an upper bound of $\beta \left\lceil \frac{\beta}{2} \right\rceil^{-1}$ on the competitive ratio of BGREEDY can be established.

We prove that this ratio is precisely $\beta \left\lceil \frac{\beta}{2} \right\rceil^{-1}$ by exhibiting a special sequence of requests. Let S_m be any graph on m vertices containing a Hamiltonian path and $H_m = \overline{S_m}$ be the complementary graph. Corollary 4.2 implies that $\lambda_2(H_m) = m$. Additionally, let L_m be the graph obtained from the complete graph on $\frac{m}{2}$ vertices, denoted $K_{m/2}$, and $H_{m/2}$ by connecting any vertex of $K_{m/2}$ to any vertex of $H_{m/2}$. By construction, $\lambda_2(L_m) = \frac{3m}{2}$.

For some bound β, let the request sequence consist of the vertices of the graph $L_{2\beta}$, such that all the β vertices of K_β are requested before the vertices of H_β. Clearly, BGreedy can only accept the first $\left\lceil \frac{\beta}{2} \right\rceil$ requests, while the optimal algorithm accepts the last β requests. This establishes the competitive ratio of BGreedy. □

Notice that BGreedy does not take into account the structure of the optimal solution on the current set of requests. The previous instance shows that unless a deterministic on-line algorithm takes into account the optimal solution on the current set of requests, it cannot achieve competitive ratio better than $\beta\left\lceil \frac{\beta}{2} \right\rceil^{-1}$. Therefore BGreedy is optimal among on-line algorithms that do not reject by choice. We next prove that no randomized on-line algorithm can achieve a competitive ratio less than $\frac{3}{2}$ against any adversary.

Lemma 4.5. *No randomized on-line algorithm for the problem of radio labeling with bound can achieve competitive ratio less than* $\frac{3}{2}$.

Proof. We prove the lower bound against the oblivious adversary; since the oblivious adversary is the least powerful one, this implies a lower bound of $\frac{3}{2}$ against any adversary. The proof consists of defining an appropriate probability distribution on the request sequences and applying Lemma 2.1.

Let $\{1, \ldots, \beta\}$ be the set of available labels. We consider the following probability distribution on the request sequences:

1. With probability p, the request sequence consists of the vertices of the graph K_β.
2. With probability $1-p$, the request sequence consists of the vertices of the graph $L_{2\beta}$, such that the vertices of K_β precede the vertices of H_β.

If the input is K_β, the optimal off-line algorithm accepts exactly $\left\lceil \frac{\beta}{2} \right\rceil$ requests, while if the input is $L_{2\beta}$, it accepts exactly β requests. Thus, the expected number of accepted requests for the optimal algorithm is

$$\mathrm{E}(OPT) = p\left\lceil \frac{\beta}{2} \right\rceil + (1-p)\beta$$

Let A be any deterministic algorithm. On input K_β, the algorithm A will accept x requests using labels from the set $\{1, \ldots, 2x-1\}$, where x is any integer number between 1 and $\left\lceil \frac{\beta}{2} \right\rceil$. The value of x is fixed after the choice of the input distribution, in order A to be the deterministic on-line algorithm maximizing the expected number of accepted requests with respect to the specific probability distribution on the request sequences, i.e. A to be the "best" deterministic algorithm. Thus, with probability p, A accepts x requests. Additionally, A has to accept exactly x requests from the graph K_β on input $L_{2\beta}$, because A is a deterministic on-line algorithm and the vertices of K_β precede the vertices of H_β. Consequently, A can accept at most $(\beta - 2x)$ requests from the subgraph H_β. Thus, the expected number of accepted requests of the best deterministic algorithm in this distribution is at most

$$\mathrm{E}(A) = px + (1-p)(\beta - x)$$

For $p = \frac{1}{2}$, this reduces to

$$\mathrm{E}(A) = \frac{\beta}{2}, \quad \text{and} \quad \mathrm{E}(OPT) = \frac{1}{2}\left(\left\lceil \frac{b}{2} \right\rceil + \beta\right) \geq \frac{3\beta}{4}$$

Thus, any deterministic on-line algorithm cannot achieve a competitive ratio less than $\frac{3}{2}$ against this probability distribution. Consequently, Lemma 2.1 implies that $\frac{3}{2}$ is a lower bound on the performance of any randomized on-line algorithm against the oblivious adversary. □

In the on-line setting, we mainly face information-theoretic questions that have to do with the value of information on the computation of a minimum cost labelling. Thus, the lower bound holds for any on-line algorithm A, and it does not depend on the running time of A. On the other hand, even if we know the entire request sequence, radio labeling is $\mathcal{NP}$-complete, and, therefore, not expected to be solvable optimally in polynomial time. However, the proof of Lemma 4.5 does not take into account the $\mathcal{NP}$-completeness of radio labelling in H_β. Consequently, a polynomial-time on-line algorithm with competitive ratio $\frac{3}{2}$ is unlikely to exist.

4.3. Radio Labeling without Bound. We continue to analyze the performance of the greedy algorithm UGreedy for radio labeling without bound, where the algorithm is not allowed to reject requests. UGreedy assigns to a newly arrived request i, $i = 1, \ldots, n$, the least integer i such that both i has not been assigned to a previous request, and the resulting radio labeling is a valid one. Since there does not exist an upper bound on the labels used, such a label i always exist.

Lemma 4.6. *The competitive ratio C_U of* UGreedy *is* $2 - \frac{3}{n} \leq C_U \leq 2 - \frac{1}{n}$.

Proof. For any sequence of n requests, the value of a labeling computed by UGreedy is at most $2n-1$, since the labeling $1, 3, \ldots, 2n-1$ is always a valid one. The upper bound on the competitive ratio follows from the fact that the value of an optimal labeling cannot be less than n.

To show the lower bound, for any even integer $n \geq 2$, let U_n be the graph on n vertices obtained from K_n by removing the edges of a Hamiltonian path. Hence, the complement of U_n is a Hamiltonian path on n vertices, and $\lambda_2(U_n) = n$. Let $u_1, u_2, \ldots, u_n$ be an ordering of the vertices of $U(n)$ according to their appearance in the Hamiltonian path contained by the complementary graph; that is, for $i = 1, \ldots n-1$, the vertices u_i and u_{i+1} are not adjacent in U_n. The adversary requests all the vertices of U_n in the following order: for $i = 0, \ldots, \frac{n}{2} - 1$, it requests the vertex u_{i+1} followed by the vertex u_{n-i}.

Claim 2. For $i = 0, \ldots, \frac{n}{2} - 2$, UGreedy assigns to the vertex u_{i+1} the label $4(i+1) - 3$ and to u_{n-i} the label $4(i+1) - 1$. Moreover, UGreedy computes a radio labelling of U_n of value $2n - 3$.

Proof. We prove the first part of the claim by induction on i. Clearly, for $i = 0$, the vertex u_1 gets the label 1, and the vertex u_n gets the label 3, since u_1 and u_n are adjacent in U_n.

Assume inductively that the claim holds for any i, $0 < i < \frac{n}{2} - 2$. Then, for $i+1$, the adversary requests the vertices u_{i+2} and $u_{n-(i+1)}$. In U_n, the vertex u_{i+2}

is adjacent to all the previously requested vertices except u_{i+1}. Since the vertex u_{n-i} has been assigned the label $4(i+1)-1$, UGREEDY assigns to u_{i+2} the label $4(i+1)+1 = 4(i+2)-3$. Similarly, since all the previously requested vertices except u_{n-i} are adjacent to the vertex $u_{n-(i+1)}$, UGREEDY assigns to $u_{n-(i+1)}$ the label $4(i+2)-1$.

Therefore, for $i = \frac{n}{2} - 2$, the UGREEDY algorithm assigns to the vertex $u_{n/2-1}$ the label $2n-7$, and to $u_{n/2+2}$ the label $2n-5$. Then, the adversary requests the vertices $u_{n/2}$ and $u_{n/2+1}$ that get the labels $2n-3$ and $2n-4$ respectively. □

Since the optimal off-line algorithm can easily compute a radio labelling of value exactly n, we obtain a lower bound of $2 - \frac{3}{n}$ on the competitive ratio of the algorithm UGREEDY. □

5. Conclusions

In this paper, we deal with some basic combinatorial problems related to Frequency Assignment in mobile and radio networks. We mainly study the approximability of these problems in the on-line setting obtaining on-line algorithms of bounded competitiveness, and proving lower bounds on the performance of any algorithm.

We remark that there exist some other cases to be considered related to on-line Frequency Assignment. One open issue is to study on-line Frequency Assignment in mobile networks under more realistic assumptions for interference. In particular, one direction is to require the frequency channels assigned at a particular base station or at adjacent base stations not only to be different but also to be far enough apart (see also [**Hale, NS**]). It has been proposed [**KK**] to use adjacent-channel constraints for the frequency channels assigned to the same base station, and co-channel constraints for the channels assigned to base stations that are at distance at most one or two from each other. However, the corresponding problems have only been studied in the framework of off-line optimization motivated by the usage of Fixed Channel Allocation (FCA) strategies. In this case, the exact number of calls per base station (demand) is assumed to be known in advance.

Another interesting issue is to define on-line radio coloring similarly to on-line radio labelling and to provide on-line algorithms and lower bounds for either general interference graphs or special cases of graphs, such as planar graphs, induced subgraphs of the triangular lattice, etc.

References

[AHLT] K. Aardal, C.A.J. Hurkens, J.K. Lenstra, S.R. Tiourine, "Algorithms for Frequency Assignment Problems", *CWI Quarterly 9*, pp. 1-9, 1996.

[AT] N. Alon, M. Tarsi, "Colorings and Orientations of Graphs", *Combinatorica* **12**(2), pp. 125-134, 1992.

[AAFLR] B. Awerbuch, Y. Azar, A. Fiat, S. Leonardi, A. Rosen, "On-line Competitive Algorithms for Call Admission in Optical Networks", Proceedings of the *4th Annual European Symposium on Algorithms*, Springer-Verlag, pp. 431-444, 1996.

[ABFR] B. Awerbuch, Y. Bartal, A. Fiat, A. Rosen, "Competitive Non-preemptive Call Control", in the Proceedings of the *5th Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 312-320, 1994.

[BLS] A. Borodin, N. Linial, M. Saks, "An optimal on-line algorithm for metrical task systems", *Journal of the Association for Computing Machinery* **39**(4), pp. 745-763, 1992.

[CM] K.M. Chandy, J. Misra, "The Drinking Philosophers Problem", *ACM Transactions on Programming Languages and Systems* **6**(4), pp. 632-646, 1984.

[CS] M. Choy, A. Singh, "Efficient Fault Tolerant Algorithms for Resource Allocation in Distributed Systems", *ACM TOPLAS* **17**(3), pp. 535-559, 1995.

[ERT] P. Erdös, A. Rubin, H. Taylor, "Choosability in graphs", Proceedings of the *West Coast Conference on Combinatorics, Graph Theory and Computing*, (Congr. Num 26), pp. 125-157, 1979.

[GJ] M.R. Garey and D.S. Johnson, *Computers and Intractability: A Guide to the Theory of NP-Completeness*, Freeman, San Francisco, 1979.

[GPT] N. Garg, M. Papantriantafilou, and Ph. Tsigas, "Distributed List Coloring: How to Dynamically Allocate Frequencies to Mobile Base Stations", Proceedings of the *8th Annual Symposium on Parallel and Distributed Processing*, pp. 18-25, IEEE Press, 1996.

[GK] S. Ghosh, M.H. Karaata, "A self stabilizing algorithm for coloring planar graphs", *Distributed Computing*, **7**, pp. 55-59, 1993.

[GL] R.A.H. Gower and R.A. Leese, "The Sensitivity of Channel Assignment to Constraint Specification", Proceedings of the *12th International Symposium on Electromagnetic Compatibility*, pp. 131-136, 1997.

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

[Hall] M. Halldórsson, "A Still Better Performance Guarantee for Approximate Graph Coloring", *Information Processing Letters* **45**, pp. 19-23, 1993.

[Hara] F. Harary, Personal communication, 1997.

[HLS] J. van den Heuvel, R.A. Leese, M.A. Shepherd, "Graph Labeling and Radio Channel Assignment", `http://www.maths.ox.ac.uk/users/gowerr/preprints.html`, 1996.

[JT] T.R. Jensen, B. Toft, "Graph Coloring problems", *Wiley-Interscience Series in Discrete Mathematics and Optimization*, 1995.

[KMS] D. Karger, R. Motwani, M. Sudan, "Approximate graph coloring by semidefinite programming", Proceedings of the *35th Annual IEEE Symposium on Foundations of Computer Science*, pp. 2-13, 1994.

[KK] S. Khanna and K. Kumaran, "On Wireless Spectrum Estimation and Generalized Graph Coloring", Proceedings of the *17th Joint Conference of IEEE Computer and communications Societies - INFOCOM*, 1998.

[LT] R. Lipton and A. Tomkins, "Online Interval Scheduling", Proceedings of the *5th Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 302-311, 1994.

[LY] C. Lund and M. Yannakakis, "On the Hardness of Approximating Minimization Problems", *Journal of the Association for Computing Machinery* **41**, pp. 960-981, 1994.

[NS] L. Narayanan and S. Shende, "Static Frequency Assignment in Cellular Networks", Proceedings of the *4th International Colloquium on Structural Information and Communication Complexity*, 1997.

[PPS] G. Pantziou, G. Pentaris, P. Spirakis, "Competitive Call control in Mobile Networks", Proceedings of the *8th International Symposium of Algorithms and Computation, ISAAC '97*, pp. 404-413, 1997.

[PY1] C.H. Papadimitriou and M. Yannakakis, "Optimization, Approximation and Complexity Classes", *Journal of Computer and System Sciences* **43**, pp. 425-440, 1991.

[PY3] C.H. Papadimitriou and M. Yannakakis, "The Traveling Salesman Problem with Distances One and Two", *Mathematics of Operations Research* **18**(1), pp. 1-11, 1993.

[PPS] R. Prakash, N. Shivaratri, M. Sighal, "Distributed Dynamic Channel Allocation for Mobile Computing", in the Proceedings of *14th ACM Symposium on Principles of Distributed Computing*, pp. 47-56, 1995.

[Rayc] A. Raychaudhuri, *Intersection assignments, T-colourings and powers of graphs*, PhD Thesis, Rutgers University, 1985.

[ST] D. Sleator and R.E. Tarjan, "Amortized Efficiency of List Update and Paging Rules", *Communications of Association for Computing Machinery* **28**, pp. 202-208, 1985.

[WR] S-W. Wang and S.S. Rappaport, "Signal-to-Interference Calculations for Balanced Channel Assignment in Cellular Communications Systems", *IEEE Transactions on Communications*, **37**, pp. 1077-1087, 1989.

[Yao] A.C. Yao, "Probabilistic Computations: Towards a Unified Measure of Complexity", Proceedings of the *18th Annual Symposium on Foundations of Computer Science*, pp. 222-227, 1977.

Computer Technology Institute, Kolokotroni 3, 26221 Patras, Greece; Department of Computer Engineering and Informatics, University of Patras, Rion, 26500 Patras, Greece

E-mail address: `fotakis@cti.gr`

Computer Technology Institute, Kolokotroni 3, 26221 Patras, Greece

E-mail address: `pantziou@cti.gr`

Computer Technology Institute, Kolokotroni 3, 26221 Patras, Greece; Department of Computer Engineering and Informatics, University of Patras, Rion, 26500 Patras, Greece

E-mail address: `pentaris@cti.gr`

Computer Technology Institute, Kolokotroni 3, 26221 Patras, Greece; Department of Computer Engineering and Informatics, University of Patras, Rion, 26500 Patras, Greece

E-mail address: `spirakis@cti.gr`