Konrad-Zuse-Zentrum
für Informationstechnik Berlin

Takustraße 7
D-14195 Berlin-Dahlem
Germany

ANDREAS EISENBLÄTTER
MARTIN GRÖTSCHEL
ARIE M.C.A. KOSTER

Frequency Planning and Ramifications of Coloring

ZIB-Report 00–47 (December 2000)

Frequency Planning and Ramifications of Coloring

Andreas Eisenblätter[†] Martin Grötschel[†] Arie M.C.A. Koster[†]

December 21, 2000

Abstract

This paper surveys frequency assignment problems coming up in planning wireless communication services. It particularly focuses on cellular mobile phone systems such as GSM, a technology that revolutionizes communication. Traditional vertex coloring provides a conceptual framework for the mathematical modeling of many frequency planning problems. This basic form, however, needs various extensions to cover technical and organizational side constraints. Among these ramifications are T-coloring and list coloring. To model all the subtleties, the techniques of integer programming have proven to be very useful.

The ability to produce good frequency plans in practice is essential for the quality of mobile phone networks. The present algorithmic solution methods employ variants of some of the traditional coloring heuristics as well as more sophisticated machinery from mathematical programming. This paper will also address this issue.

Finally, this paper discusses several practical frequency assignment problems in detail, states the associated mathematical models, and also points to public electronic libraries of frequency assignment problems from practice. The associated graphs have up to several thousand nodes and range form rather sparse to almost complete.

1 Introduction

More than a century ago, in the early 1890s, several researchers started to experiment with wireless communication via radio waves. In 1909, Marconi and Braun received the Nobel Prize in Physics for their pioneering work on the wireless telegraph. Continuing improvements of the equipment resulted in the establishment of wireless (long distance) telephony as well as radio and television broadcasting. In the last 50 years, the radio spectrum has been explored for wireless communication in many different ways. Nowadays radio waves are not only used for the already mentioned applications, but also for cellular telephone networks, radar, navigational systems, military communication, and space communication.

All these applications use frequencies in the radio spectrum to establish wireless communication. The frequencies that are applicable are limited, however. They range roughly from

[†]Konrad-Zuse-Zentrum für Informationstechnik Berlin, Takustraße 7, D-14195 Berlin, Germany.
Keywords: frequency assignment, graph coloring
Mathematics Subject Classification (MSC 2000): 05–02, 05C90, 90–02, 05C15

3 kHz to 300 GHz, corresponding to wavelengths between 100 km and 1 mm. Interference of radio signals calls for a strict management of frequency use at all levels: global, national, and regional. At the global level, the International Telecommunication Union (ITU) regulates the frequency use, whereas national agencies do the same within a country. They issue licenses to use certain frequencies for specific applications. Figure 1 gives an overview of which frequencies are currently used for which applications. For example, the most popular applications, radio, television, and cellular phone, use frequencies in the very high frequency (VHF) spectrum and ultra high frequency (UHF) spectrum. New applications such as the much discussed Universal Mobile Telecommunication System (UMTS) have to be fitted into the spectrum, whereas licenses for hardly used (ancient) applications can be retracted.

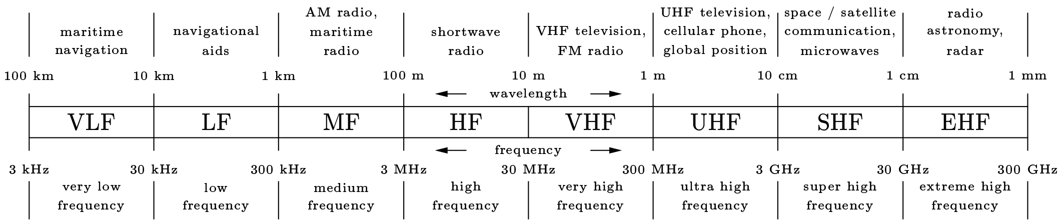

Figure 1: The radio spectrum applicable for wireless communication.

Wireless communication between two points is established with the use of a transmitter and a receiver. The transmitter emits electromagnetic oscillations. These oscillations can be modulated either via the amplitude or via the frequency itself. The receiver detects these oscillations and transforms them in either sounds or images. When two transmitters use frequencies close to one another (or to their harmonics), their signals may interfere. The level of interference depends on many aspects such as the distance between the transmitters, the geographical position of the transmitters, the power of the signals, the direction in which the signals are transmitted, and the weather conditions. In case the level of interference is high, the signal's quality may be so poor at the receiver that a proper reception is impossible. (The strength of a signal in comparison to the sum of the strengths of the interfering signals, also called noise, is expressed by a *signal-to-noise ratio.*)

The commercially usable radio spectrum is very scarce. As a consequence, frequencies are typically reused by many transmitters within one and the same network. A high performance of a network can only be achieved by carefully planning the assignment of frequencies to transmitters. The selection of the frequencies in such a way that interference is avoided or, second best, minimized, is called the Frequency Assignment Problem (FAP). The conditions to be satisfied by a frequency plan vary depending on the application. Therefore, it is not surprising that many different approaches have been suggested in the literature to solve this problem. All models, however, have one feature in common: in some way or another they are generalizations of the well known node coloring problem in an undirected graph.

In this paper, we survey the evolution of frequency assignment problems from standard vertex coloring to advanced models aiming at the minimization of the interference in a cellular radio network. Section 2 gives a glimpse on the technical side constraints that turn frequency planning into a hard combinatorial problem. The presently very important application of GSM cellular phone networks serves as our example. Next, we turn to the mathematical aspects of frequency planning. In Section 3, the development of mathematical models for frequency assignment is discussed. The theoretical and practical hardness of the problems is the topic of Section 4. We discuss the computational complexity of the presented models and characteristics of practical instances. Finally, upper and lower bounding techniques applied to the optimization problems at hand are discussed in Section 5 and 6, respectively. The paper closes with remarks on the current state of theory and practice of frequency planning.

2 GSM Cellular Mobile Phone Networks

In this section, we concentrate on terrestrial mobile cellular networks, an application that has revolutionized the telephone business in the recent years and is going to have further significant impact in the years to come. Even in this special application the frequency assignment problem has no universal mathematical model. We focus on the GSM standard (GSM stands for "General System for Mobile Communication"), which has been in use since 1992. GSM is the basis of almost all cellular phone networks in Europe. It is employed in more than 100 countries and serves several hundred million customers. The new worldwide standard UMTS (Universal Mobile Telecommunication System) is expected to become commercially available around 2002. UMTS is frequently covered in the public press at present because of the enormous amounts of money telephone companies are paying in the national frequency auctions. UMTS handles frequency reuse in an even more intricate manner than GSM: frequency or time division are used in combination with code division multiple access (CDMA) technology.

2.1 Channel Spectrum

The typical situation in GSM frequency planning is as follows. A telephone company (let us call it the *operator*) has acquired the right to use a certain spectrum of frequencies $[f_{min}, f_{max}]$ in a particular geographical region, e. g., a country. The frequency band is—depending on the technology utilized—partitioned into a set of *channels*, all with the same bandwidth Δ. The available channels are here denoted by $1, 2, \dots, N$, where $N = (f_{max} - f_{min})/\Delta$. In Germany, for instance, an operator of a mobile phone network owns about 100 channels. On each channel available, one can communicate information from a transmitter to a receiver. For bidirectional traffic a second channel is needed. In fact, if an operator acquires a spectrum $[f_{min}, f_{max}]$ he automatically obtains a paired spectrum of equal width for bidirectional communication. One of these spectra is used for mobile to base station (up-link), the other for base station to mobile (down-link) communication.

It is customary to ignore the subtle difference between channels and frequencies and to use the words as synonyms.

2.2 BTSs, TRXs, and Cells

To serve his customers an operator has to solve a number of nontrivial problems. In an initial step the geographical distribution of the communication demand for the planning period is estimated. Based on these figures, a communication infrastructure has to be installed capable to serve the anticipated demand. The devices handling the radio communication with the mobile phones of the customers are called Base Transceiver Stations (BTS). They have radio transmission and reception equipment, including antennas and all necessary signal processing capabilities. An antenna of a BTS can be omni-directional or sectorized. The typical BTS used today operates three antennas, each with an opening angle of 120 degrees. Every such antenna defines a *cell*. These cells are the basic planning units (and that is why mobile phone systems are also called *cellular phone systems*).

The capacity of a cell is defined by the number of transmitter/receiver units, called TRXs, installed for its antenna. The first TRX handles the signaling and offers capacity for up to six calls (by time division). Additional TRXs can typically handle 7 or 8 further calls—depending on the extra signaling load. No more than 12 TRXs can be installed for one antenna, i. e., the maximum capacity of a cell is in the range of 80 calls. That is why areas of heavy traffic (e. g., airports, business centers of big cities) have to be subdivided into many cells.

2.3 BSCs, MSCs, and the Core Network

In a next planning step, the operator has to locate and install the so called Base Station Controllers (BSCs). Each BTS has to be connected (in general via cable) to such a BSC, while a BSC operates several BTSs in parallel.

Every BSC, in turn, is connected to a Mobile Service Switching Center (MSC). The MSCs are connected to each other through the so called *core network*, which has to carry the "backbone traffic." The location planning for BSCs and MSCs, the design of the topology of the core network, the optimization of the link capacities, routing, failure handling, etc., constitute major tasks an operator has to address. We do not discuss here the roles of all the devices that make up a mobile phone network and their mutual interplay in detail. This brief sketch is just meant to indicate that telecommunication network planning is quite a complex task.

2.4 Channel Assignment, Hand-Over

We have seen that the TRXs are the devices that handle radio communication with the mobile phones of the customers. The operators in Germany maintain networks of about 5,000 to 15,000 TRXs and have around 100 channels available. Thus, the question arises how to best distribute the channels to the TRXs.

An operational mobile phone emits signals that allow the network to roughly keep track of which cell the mobile phone is currently "listening" to. This is done via so called control channels. If a call arrives, the network searches for the mobile phone. The device answers to the cell it was previously inactively listening to. Provided spare capacity is available, the mobile phone is assigned to one of the TRXs of the cell. If the phone moves (e. g., in a car) the communication with its current TRX may become poor. The system monitors the reception quality and may decide to use a TRX from another cell. Such a switch is called *hand-over*.

This short discussion shows that a mobile phone typically is not only in one cell. In fact, some cells must overlap, otherwise hand-overs were not possible.

2.5 Interference

Whenever two cells overlap and use the same channel, interference occurs in the area of cell intersection. Moreover, antennas may cause interference far beyond their cell limits. The computation of the level of interference is a difficult task. It depends not only on the channels, the signals' strength and direction, but also on the shape of the environment, which may strongly influence wave propagation. There are a number of theoretical methods and formulas with which interference can be quantified. Most mobile phone companies base their analysis of interference on some mathematical model taking transmitter power, distances, as well as fading and filtering factors into account. The data for these models typically come from terrain and building data bases but may also include vegetation data. Such data are combined with practical experience and extensive measurements. The result is an interference prediction model with which the so called *co-channel interference*, which occurs when two TRXs transmit on the same channel, is quantified. There may also be *adjacent-channel interference* when two TRXs operate on channels that are adjacent (i. e., one TRX operates on channel i, the other on channel $i + 1$ or $i - 1$).

Reality is still a bit more complicated than sketched before. Several TRXs (and not only two) operating on the same or adjacent channels may interfere with each other at the same time. And what really is the interference between two cells? It may be that two cells interfere only in 10% of their area but with high noise or that they interfere in 50% of their area with low noise. What if interference is high but almost no traffic is expected? How can a single "interference value" reflect such a difference in the interference behavior? There is no clear answer.

The planners have to investigate such cases in detail and have to come up with a reasonable compromise. The result, in general, is a number, the *interference value*, which is usually normalized to be between 0 and 1. This number should—to the best of the knowledge of the planners—characterize the interference between two TRXs (in terms of the model, the technological assumptions, etc., used by the operator).

2.6 Separation and Blocked Channels

If two or more TRXs are installed at the same location (or site), there are restrictions on how close their channels may be. For instance, if a TRX operates on channel i, a TRX at the same site is not allowed to operate on channels $i - 2, \dots, i + 2$. Such a restriction is called co-site *separation*. Separation requirements may even be tighter if two TRXs are not only co-site, but also serve the same cell. Separation requirements may apply also to TRXs that are in close proximity.

Moreover, due to government regulations, agreements with operators in neighboring regions, requirements from military forces, etc., an operator may not be allowed to use its whole spectrum of channels at every location. This means that, for each TRX, there may be a set of so called *blocked channels*.

2.7 A first Step into Mathematics: The Interference Graph

A feasible assignment of channels to TRXs clearly has to satisfy all separation constraints. Blocked channels also must not be used. What should one do about interference?

On our way to an adequate mathematical representation of all technical constraints, let us first introduce the *interference graph* $G = (V, E)$. G has a node for every TRX, two nodes are joined by an edge if interference occurs when the associated TRXs operate on the same channel or on adjacent channels or if a separation constraint applies to the two TRXs. With each edge $vw \in E$, two interference values, denoted by $c^{co}(vw)$ and $c^{ad}(vw)$, are associated; the number $c^{co}(vw)$ is the co-channel interference that occurs when TRXs v and w operate on the same channel, while $c^{ad}(vw)$ denotes the interference value coming up when v and w operate on adjacent channels. In general, $c^{co}(vw) \geq c^{ad}(vw)$. If a separation constraint applies to v and w then a suitable large number is allocated to $c^{co}(vw)$ and $c^{ad}(vw)$.

3 Evolution of Mathematical Models

The modeling of frequency planning problems as mathematical optimization problems has a long tradition. Metzger [44] appears to be the first to bring frequency assignment to the attention of the operations research community. Hale [27] compiled an extensive classification of the frequency planning problems of that time, phrased many of them in the graph coloring framework, and introduced notions such as T-coloring. In this section, we discuss the major generalizations of graph coloring introduced over the years in order to cover more and more aspects of frequency planning. We discuss a few theoretical results obtained for these problems. It soon turned out, though, that the various generalizations are not able to completely address all the requirements in practice. As a consequence, other modeling approaches were proposed, some of which are discussed in Subsections 3.5 and 3.6. Alternative ways to model frequency planning problems are presented in Subsection 3.7.

3.1 FAP and Coloring

One of the main difficulties in operating a wireless network is that signals within the same geographical region may interfere, resulting in a loss of quality of the received signal. In case the signal-to-noise ratio drops below a certain value (depending on the technology used), the signal may be too poor to enable communication. In particular, signals transmitted at the same frequency in the "same area" (what this means needs definition, the precise definition again depends on the technology employed) lead to substantial interference rates, the co-channel interference. To avoid these interferences, transmitters in the same area have to operate on different frequencies. Hence, neglecting all other aspects, the FAP may be reduced to coloring the vertices of a graph. Here, every vertex represents a transmitter and two vertices are connected by an edge if their transmitters are in the same area, i. e., their signals interfere if they are transmitted at the same frequency. We call this resulting graph the *conflict graph*, which clearly is a special case of the interference graph.. Colors simply correspond to frequencies.

Now, a conflict-free coloring of the vertices corresponds to an assignment of frequencies to the transmitters without co-channel interference. Note that specific channels do not play any role in this assignment, only the fact that neighboring vertices are assigned different frequencies is important. As a consequence, the minimum number of frequencies needed to obtain such an assignment is given by the chromatic number of the conflict graph G. If we define the *span* of an assignment (denoted by $sp(G)$) as the difference between maximum and minimum channel number used, then clearly the minimum span equals the chromatic number of the graph minus one (cf. Hale [27]).

3.2 FAP and T-Coloring

The signal-to-noise ratio does not only depend on the co-channel interference, but also on the adjacent-channel interference. Even farther away frequencies can influence the reception quality of the signal. In particular, frequencies that are harmonics may cause interference. So, in general not only frequencies that are equal cannot be assigned to neighboring vertices, but certain distances between the frequencies have to be obeyed. The consequences of this extension of the FAP for the vertex coloring problem are twofold. First, to measure distances, the colors have to be ordered. Assuming the channels are evenly spaced in the frequency band, the channels/colors can be numbered 1 to N. Second, a finite set $T_{vw} \subset \mathbb{Z}$ is associated with every edge $vw \in E$. This set contains the forbidden distances for frequencies f_v and f_w assigned to v and w, i. e., $f_v - f_w \notin T_{vw}$. Typically T_{vw} is symmetric, i. e., $t \in T_{vw} \Leftrightarrow -t \in T_{vw}$. In such a case T_{vw} is reduced to the non-negative elements only, and $|f_v - f_w| \notin T_{vw}$ is required. Moreover, from a practical point of view, we can assume $0 \in T_{vw}$. The (ordinary) coloring problem is the special case, where $T_{vw} = \{0\}$ holds for all $vw \in E$. This generalization of vertex coloring to practical frequency assignment is called T-coloring and was introduced by Hale [27]. For a survey, we refer to Roberts [49].

Although T-coloring is a proper generalization of vertex coloring, these two concepts have a fundamental property in common: Cozzens and Roberts [18] proved that the T-chromatic number of a graph G, $\chi_T(G)$, is equal to the chromatic number $\chi(G)$. Hence, minimizing the number of frequencies needed for a T-coloring is equivalent to minimizing the number of frequencies necessary for an assignment without co-channel interference. Minimizing the span of a T-coloring (denoted by $sp_T(G)$), however, is quite different from minimizing the span of a coloring. Figure 2 shows an example. The graph displayed in Figure 2(a) has chromatic number 3. By adding 1 to the forbidden distances T_{vw} for the edges of the cut $\delta(\{a, c\})$ (the sets T_{vw} are displayed along the edges vw in Figure 2(b)), we obtain a (more general) T-coloring problem. In this case, $sp_T(G) > \chi(G) - 1$. Moreover, no solution exists that minimizes both, the span and the number of colors, simultaneously. Every solution that uses only 3 colors has span at least 4 (for instance the coloring $f = (1, 3, 5, 1)$), whereas every solution with span 3 uses 4 colors (e. g., $(2, 4, 1, 3)$ or $(3, 1, 4, 2)$).

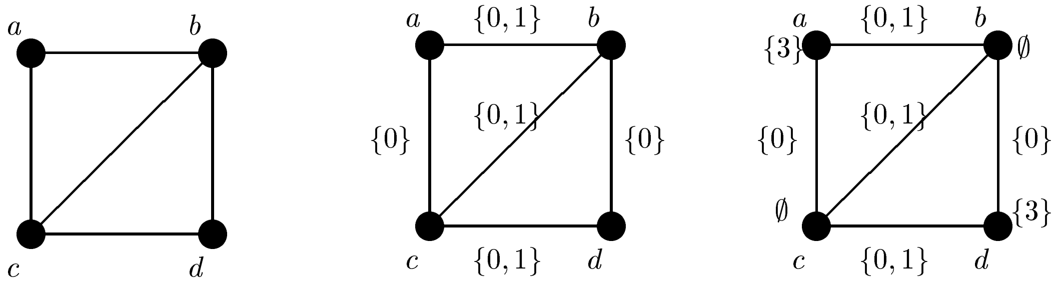

(a) coloring problem
$\chi(G) = 3,\ sp(G) = 2$

(b) T-coloring problem
$\chi_T(G) = 3,\ sp_T(G) = 3$

(c) list-T-coloring problem
$\chi_{\ell T}(G) = 3,\ sp_{\ell T}(G) = 4$

Figure 2: Example coloring, T-coloring, list-T-coloring.

The special case of T-coloring, where $T_{vw} = \{0, 1, \ldots, d(vw) - 1\}$ for some $d(vw)$ for all $vw \in E$ has received a lot of attention. The values $d(vw)$ are then called the *separation distances* and the constraints reduce to $|f - g| \geq d(vw)$ for all $vw \in E$.

3.3 FAP and List-T-Coloring

Another aspect that plays a major role in frequency planning is the (un)availability of certain channels at some transmitters. For example, GSM networks in different countries can be licensed to use the same frequencies. To avoid unforeseeable interference in border regions, network operators mutually agree not to use particular parts of their frequency bands in these areas. In general, this leads to a set of blocked frequencies B_v for every transmitter v. This generalization of coloring is known as *list-coloring*. List-coloring was first discussed by Vizing [58]; Erdős, Rubin, and Taylor [22] independently brought the problem to the attention of a wider public. Tesman [53, 55] combined list-coloring and T-coloring to *list-T-coloring*. Figure 2(c) shows an example with list-T-chromatic number $\chi_{\ell T}(G) = 3$ and list-T-span $sp_{\ell T}(G) = 4$. The above mentioned T-colorings are the only ones with span 3 and are both forbidden. Except for the sets of forbidden colors displayed at every vertex, we assume in the example of Figure 2(c) that all non-positive integers as well as all integers greater or equal to 6 are forbidden.

3.4 FAP and Set-Coloring

So far, every vertex of our graph corresponds with a single transmitter needing one frequency. In practice, transmitters are often bundled at the same location (or antenna or site, cf. Section 2). The different transmitters at the same site are usually assumed to be technically identical. Hence, the vertices in the constraint graph have the same neighborhood (except for their mutual relation), and so, combining them into a single vertex is an attractive way to reduce the graph size. In this reduced graph, multiple frequencies have to be assigned to a single vertex. The distance enforced between the frequencies assigned to the same vertex is called the *co-site* distance. In the most simple case, where only co-channel interference is considered, the problem is known as *set-coloring* (cf. [48]). Tesman [54, 55] also combined T-coloring with set-coloring resulting in *set-T-coloring*. Many results for set-coloring and set-T-coloring concern the coloring of every vertex with exactly k colors, i. e., k-tuple colorings [51].

3.5 Fixing the Span and introducing Interference Thresholds

In all problems discussed above, the objective is either to minimize the number of used frequencies or to minimize the span of the used frequencies. In practice, however, licenses for frequency bands are usually bought for long-term periods and no extension (reduction) of these bands is possible. Increasing demand for communication sooner or later leads to a shortage of frequencies, i. e., a minimum span or minimum order frequency plan needs more frequencies than are licensed. Hence, models such as list-T-coloring do not really satisfy the needs for implementable frequency plans.

One possible solution to this problem consists in fixing the span and introducing interference thresholds. The set of available frequencies F is fixed to only those that are licensed by the network operator. Let $F_v := F \setminus B_v$ be the set of available frequencies for a specific vertex. Problems such as list-T-coloring can become infeasible, because interference cannot be avoided by any frequency plan. Consider, for instance, the list-T-coloring problem of Figure 2(c). Suppose that the span is fixed to the frequencies $\{1, 2, 3, 4\}$. Since $sp_{\ell T}(G) = 4$, no feasible solution for the list-T-coloring problem exists.

Now, the question is how to evaluate such "infeasible" plans. First of all, hard and soft constraints are distinguished. Hard constraints are those frequency separation constraints and frequency blockings that have to be respected by any frequency plan, whereas violation of the soft constraints can be accepted if there is no better choice. Among the soft constraints, acceptance of violation is measured with *penalty functions*. In general, for a pair of neighboring vertices v, w defining a soft constraint, a penalty function $p_{vw} \colon F_v \times F_w \to \mathbb{R}_+$ evaluates the interference level for every pair of frequencies assigned to v and w. For our application GSM, we have hard separation constraints and soft co- and adjacent-channel constraints (cf. Sections 2.5 and 2.6). If f and g denote available frequencies, then the penalty function p_{vw} has the following format in this context:

$$p_{vw}(f, g) = \begin{cases} c^{co}(vw) & \text{if } f = g \\ c^{ad}(vw) & \text{if } |f - g| = 1 \\ 0 & \text{otherwise} \end{cases}$$

How the interference values $c^{co}(vw)$ and $c^{ad}(vw)$ are determined has been indicated in Section 2.5. A precise description goes beyond the scope of this paper. Area-based and traffic-based ratings of interference are most often used. For more on this topic we refer to [20, Section 2.3] and the references therein.

Given the interference values, mainly two solution approaches to the FAP are investigated. Frequencies are assigned to the vertices in such a way that the hard constraints are not violated and either the total penalty incurred by a solution or the maximum penalty incurred by a solution is minimized. Which of these directions is better suited for the situation depends on the intention of the operator. Consider, for example, the coloring problem (i. e., only co-channel constraints are considered) in Figure 3 and suppose $F = \{1, 2\}$ is the set of available frequencies. Penalty values are displayed at the edges. In case the total penalty is minimized, a and d are assigned the same frequency, b and c get the other. The total penalty as well as the maximum penalty is $1+\epsilon$. In case the maximum penalty is minimized, a and c are assigned the same frequency, b and d obtain the other. Now, the maximum penalty is $1-\epsilon$ but the total penalty is $2-2\epsilon$. Which approach results in the better choice?

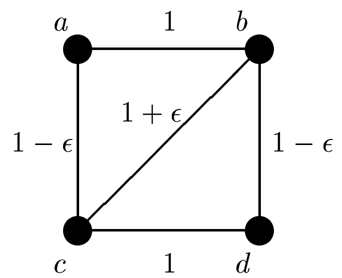

Figure 3: Example interference penalties.

In Section 3.6, we discuss the minimization of the total penalty, whereas the remaining part of this subsection deals with minimization of the maximum penalty. To minimize the maximum penalty incurred in a solution, the following approach is often applied.

Instead of computing a solution, where the maximum penalty is minimized, we search for a solution, where the incurred interference does not exceed a given threshold value P. Thus, if $p_{vw}(f, g) > P$ then the combined assignment of frequencies f and g is forbidden. This essentially reduces the problem again to a problem with only hard constraints and with the objective of finding a feasible solution, i. e., a solution in which no forbidden combinations are selected. In case such an assignment exists, a second objective is often introduced to guide the procedure to preferred feasible assignments. This second objective minimizes the number of used frequencies, the span, or the sum of all penalties below the threshold.

In case no solution exists for the given threshold value P, the maximum penalty in every assignment exceeds P. In order to find a feasible assignment, the threshold value has to be increased. If the problem restricted to the hard constraints is feasible, the threshold value P can be reduced. An assignment with minimum maximum penalty can be found by binary search on the threshold value.

3.6 Minimizing Interference

Instead of minimizing the maximum penalty, the total penalty in an assignment is often minimized. The most common way to model this so-called minimum interference frequency assignment problem (MI-FAP) is by formulating it as an integer linear program. Various formulations are proposed in the literature, see, e. g., [41, 20]. The most natural one generalizes an integer programming formulation for vertex coloring. Consider MI-FAP without hard separation constraints and in the case that $p_{vw}(f,g)$ can only take two values, 0 and p_{vw}, depending on whether $|f-g| \notin T_{vw}$ or $|f-g| \in T_{vw}$. Moreover, we assume that only one frequency has to be assigned to every vertex. The formulation uses binary variables x_{vf}, indicating whether frequency $f \in F_v$ is assigned to vertex v, and binary variables z_{vw}, indicating whether $|f-g| \in T_{vw}$ is violated. Then MI-FAP reads as

$$\begin{aligned}
\min \quad & \sum_{vw \in E} p_{vw} z_{vw} && && (1)\\
\text{s.t.} \quad & \sum_{f \in F_v} x_{vf} = 1 && \forall v \in V && (2)\\
& x_{vf} + x_{wg} \leq 1 + z_{vw} && \forall vw \in E, f \in F_v, g \in F_w \colon |f-g| \in T_{vw} && (3)\\
& x_{vf} \in \{0,1\} && \forall v \in V, f \in F_v && (4)\\
& z_{vw} \in \{0,1\} && \forall vw \in E && (5)
\end{aligned}$$

The objective (1) sums the total penalty involved in an assignment. The constraints (2) model that precisely one available channel has to be assigned to every transmitter. The inequalities (3) enforce that if $f \in F_v$ is assigned to v and $g \in F_w$ is assigned to w and $|f-g| \in T_{vw}$, then z_{vw} has to be equal to one. Since $p_{vw} > 0$, z_{vw} only equals one in an optimal solution if $|f-g| \in T_{vw}$. Finally, the constraints (4) and (5) guarantee that the variables take binary values. This formulation involves only soft constraints. Hard separation constraints can be reflected by fixing the according variable z_{vw} to zero. If all positive p_{vw}-values are equal to 1 the 0/1-program above finds a channel assignment with the least number of forbidden distances.

The formulation (1)–(5) can be extended to general penalty structures and arbitrary number of frequencies to be assigned to a vertex by introducing z_{vwfg} variables for all $vw \in E$, $f \in F_v$, $g \in F_w$ (see [41, Section 2.6]).

For the GSM frequency planning problem of Section 2, a less drastic extension of the model is sufficient. For ease of notation, let

$$\begin{aligned}
E^d &= \{vw \in E \colon d(vw) > 0\},\\
E^{co} &= \{vw \in E \colon d(vw) = 0, c^{co}(vw) > 0\}, \text{ and}\\
E^{ad} &= \{vw \in E \colon d(vw) \leq 1, c^{ad}(vw) > 0\}.
\end{aligned}$$

Instead of the variables z_{vw}, we introduce variables z^{co}_{vw} and z^{ad}_{vw} to denote violation of the co-channel and adjacent-channel constraints, respectively. Then the integer linear program reads

$$\begin{aligned}
\min \quad & \sum_{vw \in E^{co}} c^{co}(vw) z^{co}_{vw} + \sum_{vw \in E^{ad}} c^{ad}(vw) z^{ad}_{vw} && && (6)\\
\text{s.t.} \quad & \sum_{f \in F_v} x_{vf} = 1 && \forall v \in V && (7)
\end{aligned}$$

$$x_{vf} + x_{wg} \leq 1 \qquad \forall vw \in E^d, f \in F_v, g \in F_w \colon |f - g| < d(vw) \tag{8}$$

$$x_{vf} + x_{wf} \leq 1 + z^{co}_{vw} \qquad \forall vw \in E^{co}, f \in F_v \cap F_w \tag{9}$$

$$x_{vf} + x_{wg} \leq 1 + z^{ad}_{vw} \qquad \forall vw \in E^{ad}, f \in F_v, g \in F_w \colon |f - g| = 1 \tag{10}$$

$$x_{vf} \in \{0, 1\} \qquad \forall v \in V, f \in F_v \tag{11}$$

$$z^{co}_{vw} \in \{0, 1\} \qquad \forall vw \in E^{co} \tag{12}$$

$$z^{ad}_{vw} \in \{0, 1\} \qquad \forall vw \in E^{ad} \tag{13}$$

The inequalities listed under (3) in the previous IP are replaced by three sets of constraints here. In case $vw \in E^d$ and $|f - g| \in T_{vw}$, (8) enforces that at most one of the frequencies can be assigned. In case $vw \in E^{co}$, $f \in F_v \cap F_w$ can be assigned to both vertices at a penalty of c^{co}_{vw}. The constraints (9) assure that this penalty is accounted for in the objective by forcing z^{co}_{vw} to be one. Similarly, the inequalities (10) guarantee that the adjacent-channel interference is comprised in the objective if $vw \in E^{ad}$ and the frequencies f and g differ by one.

3.7 Alternatives

Besides minimizing the total interference or the maximum interference level, other models for practical frequency planning have been proposed in the literature. Most of these models deal with partial assignments. So far, we have assumed that a complete assignment with some interference is prioritized above assignments that are interference free but do not satisfy all requests for channels. This priority for complete assignments is due to the network operator's wish to offer wireless communication to as many customers as possible, tolerating some degradation in quality.

An alternative can be to retain the quality of the offered connections, but to compromise on the availability of the service. This means that we assign fewer frequencies than requested to some antennas, but that the interference in the resulting network stays below a required signal-to-noise ratio. The objective in such an approach is to maximize the service availability or equivalently minimizing the network blocking probability. In both, Chang and Kim [16] and Mathar and Mattfeldt [43], formulas for the blocking probability of a vertex are derived, resulting in a nonlinear objective function. As in the case of minimizing the interference, formulating the problem as an integer program is the common way to handle these problems. Since nonlinear integer programs are out of reach for the current algorithmic machinery, the nonlinear objective function is linearized or a simplified objective is selected that minimizes the number of requested but unassigned frequencies (for more information, see [41, Section 2.5]).

The interference in all models considered so far is based on the relation of two transmitters. However, the signal-to-noise ratio depends on all transmitters in the same region simultaneously. Although the binary relations between the transmitters restricts the interference levels, it cannot avoid an incidental exception of the overall interference due to this effect. Fischetti et al. [23] show that bounding the overall interference level at a receiver can be modeled as a linear inequality that can be added to the formulation (6)–(13).

4 Theoretical and Practical Hardness

In this section, we discuss the hardness of frequency assignment problems from a theoretical as well as from a practical point of view. We summarize results on the computational complexity and present the main characteristics of several sets of benchmark problems.

4.1 Computational Complexity

In view of the $\mathcal{NP}$-hardness of vertex coloring [39], it is not surprising that all presented versions of frequency assignment are $\mathcal{NP}$-hard as well. There are stronger results. For instance, list coloring is $\mathcal{NP}$-complete even for special graphs for which the vertex coloring problem can be solved in linear time, e. g., for interval graphs [10]. Also the (negative) results on the approximation of vertex coloring [8] have their direct consequences for frequency assignment.

Besides vertex coloring, many other $\mathcal{NP}$-complete problems are closely related to frequency assignment: maximum clique, minimum edge-deletion k-partition, maximum k-colorable subgraph, maximum k-cut, minimum k-clustering sum, and maximum frequency allocation, to name a few. For each of these problems, computational complexity results can be transfered to various FAPs. From the results for minimum edge-deletion k-partition, collected in Ausiello et al. [7] (see also [20, Section 3.2]), for instance, we can derive that

- deciding whether an instance of MI-FAP allows a feasible assignment is $\mathcal{NP}$-complete,
- MI-FAP is strongly $\mathcal{NP}$-hard,
- unless $\mathcal{P}$=$\mathcal{NP}$, MI-FAP is not in $\mathcal{APX}$, and
- unless $\mathcal{P}$=$\mathcal{NP}$, an approximation of MI-FAP within $\mathcal{O}(|E|)$ for $|F| \geq 3$ is impossible in polynomial time.

The last two results directly follow from the $\mathcal{NP}$-hardness of finding a feasible solution. However, even in the case that a feasible assignment can be found easily (e. g., no separation constraints), these results are valid. This shows that the complexity of solving MI-FAP is not solely due to the feasibility problem.

4.2 Practical Instances

Even though all variants of FAP are theoretically hard, instances arising in practice might be either small or highly structured such that enumerative techniques, such as branch-and-bound or special purpose methods, are able to handle these instances efficiently. This is typically not the case. Frequency assignment problems are also hard in practice in the sense that nobody is able to routinely solve large, practically relevant instances to optimality or with a good quality guarantee.

In many areas of optimization most of the companies using optimization techniques are very reluctant to publish data of real instances. They even refuse to slightly modify the data, fearing that competitors may, nevertheless, gain insight into what they are doing. The situation is much more positive in the FAP area. Several sets of benchmarks for frequency assignment are available. A compilation of many of these benchmarks can be found at `http://fap.zib.de`, the FAP web site. The benchmark instances originate from different applications and therefore also involve different models. For each set, we discuss their origin, the frequency assignment variant to be dealt with, and the most important results.

Philadelphia benchmarks: The Philadelphia instances were among the first discussed in the literature [5]. The original instance and certain variants have been widely used to substantiate algorithms and lower bounds for the Minimum Span FAP (MS-FAP). The Philadelphia instances are characterized by 21 hexagons, representing the cells of a cellular phone network around Philadelphia (see Figure 4). Until recently, it was common practice to model wireless phone networks as hexagonal cell systems. For each cell, a demand for frequencies is given. Figure 4(b) shows the demand (in channels) for the original instance P1. The demand vectors of the other instances, in conformity with [57] denoted by P2–P9, can be found at FAP web [21]. The objective is to find a set-T-coloring with minimum span. In the basic model, interference of cells is characterized by a co-channel reuse distance d. No interference occurs if and only if the centers of two cells have mutual distance $\geq d$. In case the mutual distance is less than d (normalized by the radius of the cells), it is not allowed to assign the same frequency to both cells. This pure co-channel case is generalized by replacing the reuse distance d by a series of non-increasing values $d^0, \ldots, d^k$ and corresponding forbidden sets $T^0 \subseteq \ldots \subseteq T^k$. The following relation holds:

$$T_{vw} = T^{j-1} \text{ whenever } d^j \leq d_{vw} < d^{j-1}, j \in \{1, \ldots, k\},$$

where d_{vw} is the distance between the cell centers. For the Philadelphia instances, the sets T^j are taken as $T^j = \{0, \ldots, j\}$. For instance P1, the values $d^0, \ldots, d^5$ are $2\sqrt{3}, \sqrt{3}, 1, 1, 1, 0$. So, frequencies assigned to the same site should be separated by at least 4 other frequencies, whereas frequencies assigned to neighboring sites should be at a distance of at least 2, and frequencies assigned to a second and third "ring" of cells should still differ. For the other instances, the reuse distances can be found in [21], where the results presented in the literature are also summarized. By now, all instances are solved to optimality.

CALMA benchmarks: The military usage of field phones also leads to FAPs. These FAPs have the property that each connection consists of two movable phones. To each connection we have to assign two frequencies at a fixed distance of each other. Thus, all frequencies can be given as combinations of two with this fixed distance. Instances are available in the EUCLID CALMA (Combinatorial ALgorithms for Military Applications) project. In the CALMA project (1993–1995) researchers from England, France, and the Netherlands tested different combinatorial algorithms on the same set of frequency assignment problems. The set contains minimum order problems, as well as minimum span and minimum interference problems with a predefined set of frequencies for every transmitter. Neglecting the fixed distance constraints for the moment, the minimum order and minimum span problems are, in fact, list-T-coloring problems. Eleven real-life instances were provided by CELAR (Centre d'ELectronique

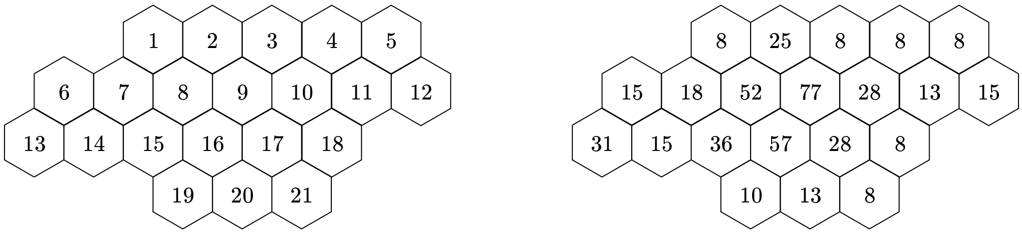

(a) network structure (b) frequency demand of instance `P1`

Figure 4: Philadelphia benchmark instances.

de l'ARmement, France). A second set of 14 instances was made available by the research group of Delft University of Technology. These GRAPH (Generating Radio Link Frequency Assignment Problems Heuristically) instances were randomly generated by Van Benthem [9] and have the same characteristics as the CELAR instances.

In addition to minimum separation constraints, the instances also contain equality constraints. They are used to model that two frequencies at a fixed distance have to be assigned to the corresponding transmitters/vertices. The distance is the same for all constraints, and every vertex is contained in exactly one equality constraint. Moreover, to every available frequency there exists only one *matching* frequency at the fixed distance. So, the problem may be rephrased to assigning pairs of frequencies to pairs of vertices. This reduces the problem size by a factor of 2. Up to four of the original distance constraints have now to be considered between two pairs of vertices. Hence, the T-coloring aspect is lost in this way.

The number of available frequencies for a vertex is 40 on average. The number of vertices ranges from 200 to 916, whereas the number of edges in the conflict graph is in between 1200 and 5500. We have to assign one frequency to every vertex. The instances as well as the results can be found in [21] (see also [3]). All minimum span instances can be solved quite efficiently. Also, all but one of the minimum order instances are solved to optimality. For the MI-FAP instances, the results are more diverse. At the end of the project, only upper bounds were available. By now, 7 out of the 11 instances are solved. For the remaining ones, lower bounds have been found differing between 57.3% and 98.2% from the best known solution value.

COST 259 benchmarks: COST (COperation européenne dans le domaine de la recherche Scientifique et Technique) is a European Union Forum for cooperative scientific research. The COST 259 project on *Wireless Flexible Personalized Communications* ran from the end of 1996 to the beginning of 2000. Working groups on Radio System Aspects, Antennas and Propagation, and Network Aspects were formed and dealt with different aspects of mobile radio communications. The subgroup SWG 3.1 of the Network Aspects working group compiled a library of GSM frequency planning scenarios. The intention was to allow comparisons of available frequency planning methods as well as to stimulate and to support the development of new methods.

In total, 32 realistic GSM network planning scenarios were collected and served as benchmark for comparing algorithms within COST 259. The scenarios together with several contributed frequency plans are made available at [21]. The primary objective is to minimize the total interference. The solutions and the lower bounds contributed so far leave space for improvements.

Tables 1 and 2 show various characteristics of selected benchmark instances. Besides B2, several other B[d] instances exist, differing in the average and maximum number of TRXs or in the propagation model, for example. Instance B2 is, in fact, the instance BRADFORD_NT-2-EPLUS of the COST 259 project. Table 1 contains figures concerning the underlying GSM network, namely: the *number of sites* in the planning area; the *number of cells* (no site hosts more than three cells); the *average number of TRXs per cell*; the *maximum number of TRXs per cell*; the *spectrum* size or, in the presence of globally blocked channels, the sizes of the contiguous portions in the spectrum; the *minimum number of channels* available in a cell; and the *average number of non-blocked channels* in a cell. Notice that in each of the scenarios with globally blocked channels, the resulting gap in the spectrum exceeds the maximal required separation. Thus, there is no direct coupling between distinct contiguous portions of the spectrum.

	# sites	# cells	avg. TRXs/cell	max. TRXs/cell	spectrum size	min. # channels	avg. # channels	avg. # adj. cells	max. # adj. cells	diameter
B2	649	1886	1.17	5	75	75	75.00	256.38	779	5
SIE1	179	506	1.84	4	20, 23	43	43.00	41.76	104	7
SIE2	86	254	3.85	6	4, 72	76	76.00	121.98	225	3
SIE3	366	894	1.82	3	55	12	51.25	75.76	267	10
SIE4	276	760	3.66	5	39	39	39.00	70.31	184	5

Table 1: Scenario characteristics COST 259 benchmarks.

Three further figures are given for each scenario in Table 1. These are easiest explained in terms of the carrier network that is obtained if each cell operates only one TRX. Then, every carrier in the network corresponds to a cell, and the graph underlying this network reflects the relations among the cells. For this graph, the average degree *(average number of adjacent cells)*, the maximum degree *(maximum number of adjacent cells)*, and the *diameter* (of the largest connected component) are listed.

We point out a few peculiarities of those graphs. Looking at the column showing the average number of adjacent cells, we see that in general a cell is adjacent to surprisingly many other cells. The maximum degree column reveals that some cells may be adjacent to more than half of the other cells. The last column shows that the cells are rather "nearby" in almost all scenarios (recall that the diameter in a connected graph is the longest shortest path between two vertices).

More graph-related figures for the carrier networks are displayed in Table 2, i.e., the network, where every TRX is represented by a vertex. In the first column, the label of the associated scenario is given. The next block contains information on the graph, namely, its number of vertices, its edge density (that is, number of edges relative

	$\lvert V\rvert$	density [%]	avg. degree	max. degree	max. clique	$\lvert\{e \in E : d(e) \neq 0\}\rvert$	$\lvert\{e \in E : c^{co}(e) \neq 0\}\rvert$	$\lvert\{e \in E : c^{ad}(e) \neq 0\}\rvert$
B2	2214	13.50	299.0	916	93	10284	320684	6871
SIE1	930	9.03	84.0	209	52	6039	33002	9911
SIE2	977	49.17	480.4	877	182	17761	216912	25615
SIE3	1623	9.18	149.1	519	78	23093	97861	15069
SIE4	2785	10.50	292.3	752	100	27964	379052	26445

Table 2: Characteristics of carrier networks COST 259 benchmarks.

to the maximal possible number of $\frac{|V|\,(|V|-1)}{2}$), the average degree of its vertices as well as the maximal degree, and the size of a maximum clique. The third block addresses the minimum separation requirements as specified by the labeling d, the co-channel interference labeling c^{co}, and the adjacent-channel interference labeling c^{ad}, respectively. For all three labelings, we list the total number of edges involved. Table 2 shows that in all the carrier networks, the average degree of a carrier is significantly higher than the number of available channels. This is an indication that a great deal of coordination among the channel assignments for different cells is necessary in order to produce good frequency plans. Moreover, the size of the maximum cliques in most carrier networks is larger than the number of available channels. Hence, every feasible assignment for those carrier networks must incur interference.

ROADEF'2001 challenge: This challenge is a follow-up of the CALMA project. It considers frequency assignment problems with polarization constraints, i. e., in addition to the assignment of frequencies, also a polarization direction, horizontal or vertical, have to be assigned. The required frequency separation distances depend on the polarization. The data set has again been made available by CELAR. More information can be found at the ROADEF'2001 web-site [47].

Miscellaneous: Several other instances are used in the literature to test algorithms for solving frequency assignment problems. We refer to [41, Section 2.2.4] for an overview of the instances and their application. An instance that is not mentioned there, is the CNET France Telecom benchmark provided by Caminada [15]. It contains raw data of a GSM network such as details of base station locations and propagation information.

5 Heuristic Approaches

The $\mathcal{NP}$-hardness of both, solving an FAP and finding solutions that are guaranteed to be close to optimal, justifies the development of (fast) methods which neither certainly produce close to optimal assignments nor even feasible assignments. These heuristic approaches are the topic of this section.

Analogous to Section 3, this section is broken into subsections for the various models. We start with a discussion of some coloring heuristics, followed by heuristics for T-coloring and list-coloring. Finally, state-of-the-art methods for interference minimization are considered. For a detailed survey of heuristics for frequency assignment we refer to [2].

5.1 Coloring Heuristics

Heuristics for the vertex coloring problem have been studied extensively for years. Such heuristics can be applied to FAPs restricted to vertex coloring, but they can also be adjusted for more advanced problems such as T-coloring or list-coloring. Many heuristics have a greedy nature. The two basic types are *sequential coloring* and *maximal stable set coloring.* In both cases, the vertices are ordered in advance. In the sequential coloring approach, each vertex is colored with the least color not conflicting with the previously colored vertices according to the given vertex order. In maximal stable set coloring, one color at a time is assigned to as many vertices as possible according to the vertex order, i. e., for every color maximal stable sets are greedily constructed. Popular orderings are *lowest degree last* or *highest degree first.* In [35], 13 variants of these polynomial graph coloring algorithms are compared. Without exception, there exist infinite classes of graphs for which each of the algorithms performs poorly. On the other hand, some versions of these heuristics have been proven to produce optimal assignments for special cases in polynomial running time. These cases, however, do not seem to appear in the practice of frequency assignment. We refer to Toft [56] and the references therein for more information.

Another heuristic that has been applied in the context of frequency planning problems is the DSATUR algorithm developed by Brélaz [13]. It uses the *saturation degree* of a vertex: the number of different colors assigned to its neighbors. Repeatedly, the vertex with highest saturation degree is selected, ties are broken by highest degree, and the least possible color is assigned to this vertex. The major difference between this algorithm and the greedy algorithms described before is that the order is not static but depends on the previously assigned vertices.

5.2 Heuristics for Minimizing the Span

The first heuristics for minimizing the span have been proposed as early as the 1970s for the Philadelphia instances, see [12, 60]. Similar to the sequential coloring approach to vertex coloring, the frequencies are assigned to the vertices according to some order of the vertices. In Sivarajan, McEliece, and Ketchum [50] several variants of this greedy algorithm are tested on 13 Philadelphia instances.

More generally, Cozzens and Roberts [18] extended the greedy sequential coloring algorithm to T-coloring. Tesman [53] generalized this algorithm to k-tuple-T-coloring. For special classes of graphs, these algorithms provide the optimal span (cf. [45] for a survey). Wang and Rutherforth [59] use a greedy algorithm as starting heuristic only, then a local search procedure is applied, where the assignment of two vertices is interchanged as long as the objective function value improves.

The DSATUR algorithm has been extended to T-coloring by Costa [17]. More elaborate, but also more time-consuming heuristics for minimizing the span such as simulated annealing, tabu search, and genetic algorithms are described in [31].

5.3 Heuristics for Interference Minimization

For interference minimization, a wide variety of heuristics has been applied. Like for minimizing the span, we distinguish fast heuristics from more elaborate, but rather time-consuming ones.

In [11], a variant of the DSATUR heuristic is used for minimum interference GSM frequency planning. Roughly speaking, the *saturation degree* now keeps track of how many channels from the spectrum are no longer available for each of the remaining unassigned carriers. The *spacing degree* sums the separation distances to unassigned neighbors. So, this degree represents how much impact assigning all of a vertex' still unassigned neighbors would have on its own assignability. We repeatedly select the vertex for which the saturation degree is maximal among those for which the spacing degree is maximal. In this way a vertex for which the assignment has a large impact on its neighbors is handled before its neighbors are. Moreover, vertices with high saturation degree are assigned as soon as possible. Another extension of DSATUR, called DSATUR with costs in [11], deals with the marginal cost of assigning a frequency to a vertex in a partial assignment. The vertices are sorted according to these costs and the frequency of minimum cost is assigned repeatedly.

Besides these constructive heuristics, improvement heuristics have been applied. For the GSM frequency planning benchmarks of the COST 259 project, the best known algorithm is due to Hellebrandt and Heller [28]. They propose to use the *Threshold Accepting* paradigm proposed in [19]. Threshold Accepting is a variant of Simulated Annealing with a deterministic acceptance criterion: a proposed change of the solution, also called a move, is accepted if an improvement is achieved or the deterioration is below a threshold. The value of this threshold declines with the progress of the algorithm. The basic scheme is augmented here by alternating between random changes and local optimization. The cost of a move is the difference between the cost of the assignments after and before the move.

Figure 5 shows an example of a carrier network, where DSATUR followed by a very simple local search heuristic (1-opt), improved the solution provided by a network operator by 96%. In this plot, sites are represented by dots in the Euclidean plane according to their geographical coordinates. Two sites are connected by a colored line if the frequency assignments results in interference among the TRXs from the two sites. If the line is drawn in a pale color, then the interference is small. With increasing interference, the color of the line turns into black.

Many other techniques, from genetic algorithms to neural networks, have been applied to MI-FAP. For the CALMA minimum interference benchmarks, for example, a genetic algorithm with optimized crossover turned out to produce the best results [40].

The running times of the heuristics differ substantially. The fastest among the methods computes an assignment within a few seconds to minutes, whereas the more complex methods have running times up to a few days on realistic planning scenarios. The tendency in practice is to use fast methods within the interactive planning cycle; the more complex methods are sometimes used to generate the final frequency plan.

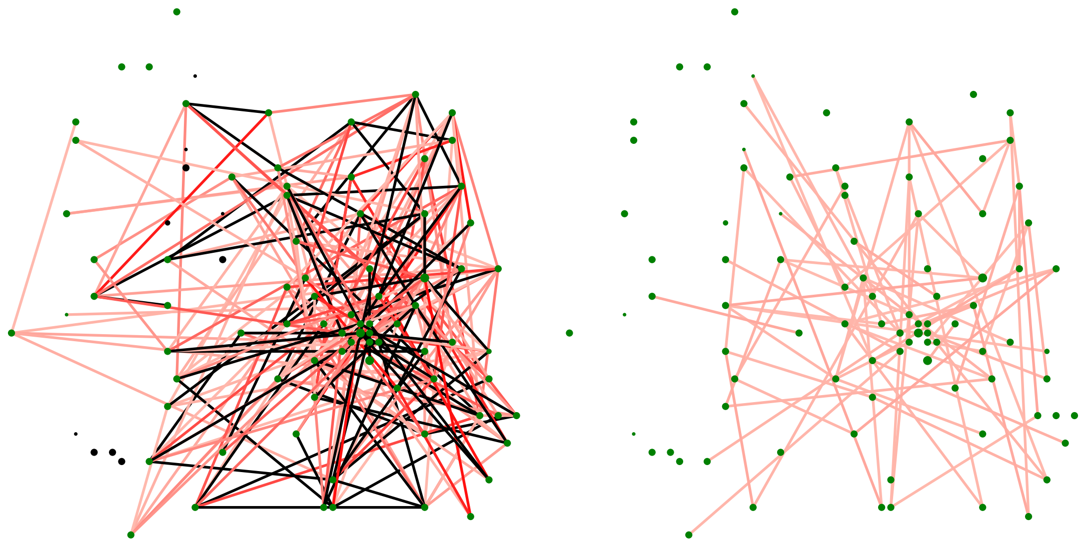

(a) operated frequency plan (b) optimized frequency plan

Figure 5: Interference reduction of 96% by heuristics.

6 Lower Bounds

As outlined in Section 5, various heuristics seem to be capable of solving frequency assignment problems in a satisfactory way—at least compared to some of the traditional planning methods employed in practice. From a mathematical point of view, a comparison with results of someone else or some other heuristic is insufficient. One would like to know better! Is it possible to give a provable quality guarantee for a heuristic solution? In principle, this is of course doable. However, coming up with acceptable quality guarantees for frequency plans appears to be quite difficult (both in theory and practice); further ideas are needed badly in this area!

All the optimization problems associated with frequency planning discussed in this paper are minimization problems. Heuristics find feasible solutions for such problems and the best of these solutions is hoped to yield a good upper bound for the objective function value of a minimum solution. To substantiate such a hope, a good lower bound for the optimum value must be found. We show in this section how lower bounds for some of the versions of the frequency assignment problems can be obtained. We also indicate how big the gap between upper and lower bound can be, despite serious efforts.

6.1 LP Lower Bounds for Coloring

We begin with describing the LP approach to combinatorial optimization, one of the most successful methods for the solution of integer programs. The idea is simple. We formulate the combinatorial problem as an integer program, drop the integrality constraint and solve the resulting linear program. One hopes that the optimum LP value is a good lower bound for the optimum value of the combinatorial problem.

Let us describe this approach for the weighted clique problem (finding a set of pairwise adjacent nodes in a node-weighted graph with maximum total weight). Let a graph $G = (V, E)$ with nonnegative weights $c_v, v \in V$, be given. Observing that a clique and a stable set intersect in at most one node, it is trivial to see that the integral feasible solutions of the following LP

$$(C) \qquad \begin{array}{lll} \max & c^T x & \\ & \sum_{v \in S} x_v \le 1 & \forall S \subseteq V, S \text{ stable} \\ & x_v \ge 0 & \forall v \in V \end{array}$$

are exactly the incidence vectors of the cliques of G. Denoting by $\omega(G, c)$ the maximum weight of a clique in G and by $\mathcal{S}$ the set of all stable sets in G we have

$$\begin{aligned} \omega(G,c) &= \max\Big\{c^T x \colon \sum_{v \in S} x_v \le 1 \,\forall S \in \mathcal{S}, x \ge 0, x \text{ integral}\Big\} \\ &\le \max\Big\{c^T x \colon \sum_{v \in S} x_v \le 1 \,\forall S \in \mathcal{S}, x \ge 0\Big\} =: \omega'(G,c) \end{aligned}$$

The LP value $\omega'(G, c)$ is sometimes called the *fractional (weighted) clique number*. Employing the duality theorem of linear programming we obtain

$$\omega'(G,c) = \min\Big\{\sum_{S \in \mathcal{S}} y_S \colon \sum_{S \ni v} y_S \ge c_v \,\forall v \in V, y_S \ge 0 \,\forall S \in \mathcal{S}\Big\}$$

Observing that, with each $S \in \mathcal{S}$, each subset of S also belongs to $\mathcal{S}$, it is clear that we can replace the inequality sign in $\sum_{S \ni v} y_S \ge c_v$ by an equality sign. Adding integrality we can continue the chain of equations and inequalities

$$\begin{aligned} \omega'(G,c) &= \min\Big\{\sum_{S \in \mathcal{S}} y_S \colon \sum_{S \ni v} y_S = c_v \,\forall v \in V, y_S \ge 0 \,\forall S \in \mathcal{S}\Big\} \\ &\le \min\Big\{\sum_{S \in \mathcal{S}} y_S \colon \sum_{S \ni v} y_S = c_v \,\forall v \in V, y_S \ge 0 \,\forall S \in \mathcal{S}, y_S \text{ integral } \forall S \in \mathcal{S}\Big\} \\ &=: \chi(G,c) \end{aligned}$$

The last minimization problem is nothing but an integer programming formulation of the weighted coloring problem: we seek a collection of stable sets (the integral value y_S is the multiplicity of the stable set $\mathcal{S}$) such that each node $v \in V$ is contained in exactly c_v sets. The ordinary coloring problem is the special case where $c_v = 1$ for all $v \in V$, i.e., $\chi(G) = \chi(G, \mathbb{1})$.

The process described above is a natural way to view the coloring problem as an IP. Solving its LP relaxation amounts to solving a linear program with $|V|$ equations and an exponential number of nonnegative variables $y_S, S \in \mathcal{S}$. Instead of that, one can, equivalently, solve its dual problem (in $|V|$ nonnegative variables and one constraint for each stable set). The large number of constraints (or variables) is not a principal difficulty (see [26] for a detailed account of this topic). However, it was shown in [25] that solving this LP is $\mathcal{NP}$-hard, i. e., computing $\omega'(G, \mathbb{1})$ in order to obtain a lower bound on $\chi(G, \mathbb{1})$ cannot be achieved in polynomial time, unless $\mathcal{P}$=$\mathcal{NP}$. This is a real pity since $\omega'(G, \mathbb{1})$ appears to be a quite good lower bound for $\chi(G, \mathbb{1})$ while, surprisingly, $\omega'(G, \mathbb{1})$ is a poor upper bound for $\omega(G, \mathbb{1})$, in general. We demonstrate this fact by means of a probabilistic argument.

Let us consider a random graph $G_{\frac{1}{2}}$, i. e., $G_{\frac{1}{2}}$ has $n = |V|$ nodes and, for each pair u, v of nodes, we throw a coin to decide whether the edge uv belongs to G or not. A lot is known about $G_{\frac{1}{2}}$, see [38] for a survey. For example, $\alpha(G_{\frac{1}{2}}) = \omega(G_{\frac{1}{2}}) = 2\log_2 n$ and $\chi(G_{\frac{1}{2}}) = \frac{n}{2\log_2 n}$ holds almost surely. We can use this fact to compute the expected value of $\omega'(G_{\frac{1}{2}})$. Namely, if we set $\overline{x}_v := \frac{1}{2\log_2 n}$ for all $v \in V$ then, since each stable set S satisfies $|S| \leq 2\log_2 n$ almost surely the vector $\overline{x} = (\overline{x}_v)_{v \in V} \in \mathbb{R}^n$ is almost surely feasible for (C), and thus, $\mathbb{1}^T\overline{x} = \frac{n}{2\log_2 n}$ is almost surely a lower bound on the optimum $\omega'(G_{\frac{1}{2}}, \mathbb{1})$ of (C). Since $\chi(G_{\frac{1}{2}})$ is almost surely $\frac{n}{2\log_2 n}$ and $\chi(G_{\frac{1}{2}})$ is an upper bound on $\omega'(G_{\frac{1}{2}}, \mathbb{1})$ we know that $\omega'(G_{\frac{1}{2}}, \mathbb{1})$ is almost surely equal to $\chi(G_{\frac{1}{2}})$. It is interesting to observe that, for the random graph $G_{\frac{1}{2}}$, $\omega'(G_{\frac{1}{2}}, \mathbb{1})$, the value of the LP relaxation of $\omega(G_{\frac{1}{2}})$, is far away from $\omega(G_{\frac{1}{2}})$ but that it provides an excellent lower bound for $\chi(G_{\frac{1}{2}})$. The trouble, though, is that except for special cases such as perfect graphs (see [26]), we do not know how to compute $\omega'(G_{\frac{1}{2}}, \mathbb{1})$ efficiently.

6.2 A TSP Lower Bound for T-Coloring

For T-coloring and its variants, many (graph-)theoretical results have been derived over the years. The survey of Murphey, Pardalos, and Resende [45] discusses these results in depth. In this subsection, we focus on one of the most interesting ideas to derive lower bounds for minimum span T-coloring.

Consider a T-coloring problem with separation distances $d(vw)$ for all $vw \in E$, i. e., $T_{vw} = \{0, \ldots, d(vw) - 1\}$. A feasible solution of this T-coloring problem is an assignment $f\colon V \mapsto \mathbb{Z}_+$ such that $|f(v) - f(w)| \geq d(vw)$ for all $vw \in E$. Let us assume that f is a feasible T-coloring, and let π be a permutation of the vertices in V such that $\pi(v) < \pi(w)$ if $f(v) < f(w)$. Hence, vertices with a larger color succeed those with a smaller color in the ordering, whereas the ordering of vertices with identical colors is arbitrary.

Let us define a complete graph G' on the vertices V of G with edge weights $d'(vw) = d(vw)$ if $vw \in E$ and $d'(vw) = 0$ otherwise. The order π defines a Hamiltonian path in G', and if we assume that f has span k, then the Hamiltonian path has weight at most k. In this fashion, every feasible T-coloring of G gives rise to one (or more) Hamiltonian path(s) in G', and in each case the span of the T-coloring is bounded from below by the length of the associated Hamiltonian path. Hence, a Hamiltonian path in G' of minimum weight provides a lower bound on the span of any T-coloring.

Finding a Hamiltonian path of minimum weight is traditionally done by solving a Traveling

Salesman Problem (TSP). Therefore, the lower bound derived this way is known as the TSP bound. This relation between T-coloring and the Hamiltonian path problem has first been observed by Raychaudhuri [46] and is later used by Roberts [49] and Janssen and Kilakos [32, 33]. The TSP as well as many of its variations have received considerable attention in the literature (see, e. g., [36] for a survey and [6] for recent progress). TSPs are typically solved using integer programming techniques. Note that also a lower bound on the minimum weight of a Hamiltonian path provides a lower bound on the minimum span of a T-coloring.

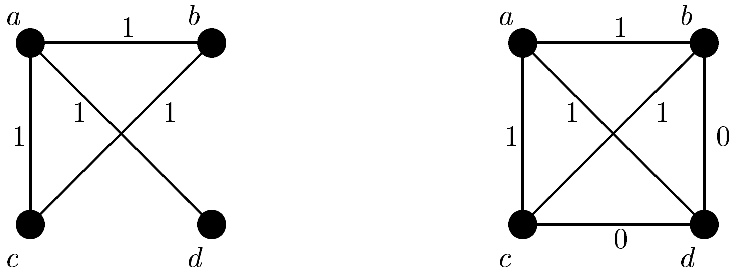

(a) Graph G (b) Graph G'

Figure 6: Example TSP lower bound.

The TSP lower bound for the minimum span of a T-coloring has a remarkable (nasty) property. Induced subgraphs of G may provide better TSP bounds than the TSP bound of G. Consider the example in Figure 6(a). The values at the edges denote the separation distances $d(vw)$. The corresponding graph G' with weights $d'(vw)$ is displayed in Figure 6(b). A Hamiltonian path of minimum weight in G' is given by a, b, d, c with weight 1. Hence, the TSP bound is 1. If we, however, restrict us to the subgraph induced by a, b, c, a TSP bound of 2 is obtained. So, the TSP bound might be improved by restricting the bound computation to a "dense part" of the conflict graph. In [33], this technique is used to prove the optimal value of Philadelphia instance P1, for example.

Recently, Allen, Smith, and Hurley [4] proposed another way to improve the TSP bound. Consider again the example in Figure 6. On the Hamiltonian path a, b, d, c, the length of the subpath b, d, c is 0, although a separation distance of $d'(bc) = 1$ exists. So, the bound can be improved by forcing that every path between b and c has weight at least $d'(bc)$. In [4], the TSP integer programming formulation is augmented by introducing so-called excess-variables in order to guarantee this property of the minimum weighted path.

6.3 The LP Lower Bound for Minimizing Interference

For the MI-FAP, lower bounds can be derived by the linear programming relaxation of (1)–(5). In the linear programming relaxation, the integrality constraints (4) and (5) are replaced by $x_{vf} \geq 0$ and $z_{vw} \geq 0$, respectively. This relaxation provides a lower bound on the optimal value. Computational experiments with this lower bound, however, did not yield satisfying results (cf. [1]). Also, the linear relaxation of other formulations (cf. [34, 41, 20]) did not provide strong lower bounds. The poor quality of the LP bound is mainly due to the freedom to select frequencies fractionally. In this way, interference can often be avoided almost completely in a fractional solution.

The branch-and-cut approach for integer programming could not be successfully applied to MI-FAPs either. The major obstacle seems to be that for both, the linear relaxation and the integer program many solutions with (almost) identical objective function value exist. In fact, there often exist permutations of a solution with the same objective function value. These symmetries lead to many equivalent branches in a branch-and-cut approach. Moreover, the addition of cutting planes merely drives the solution to another solution with same objective value.

An exception from these negative results is obtained for the CALMA MI-FAP instances. In [42] (see also [41]), integer programming is successfully applied after aggregation of frequencies to obtain lower bounds. As an alternative to integer programming, a dynamic programming algorithm based on a tree decomposition of the conflict graph is applied. Not only lower bounds could be derived using this approach, but also some of the CALMA problems could be solved to optimality (see [41] for more information). This method heavily depends on the specific structure of CALMA frequency assignment, and this clearly limits the applicability to other frequency planning problems.

6.4 A Semidefinite Lower Bound

Semidefinite programming is the task of minimizing a linear objective function over the convex cone of positive semidefinite matrices subject to additional linear inequalities. Significant research efforts in the recent years have revealed that many of the algorithmic approaches to linear programming can be tailored to work for semidefinite programs efficiently as well, see Helmberg [30]. In this subsection, we present a lower bound for the GSM frequency planning problem based on semidefinite programming.

Let us modify (in fact, simplify) the frequency planning problem by

- dropping the adjacent-channel interference (i. e., by setting $E^{ad} := \emptyset$),
- dropping the local blockings (setting $B_v := \emptyset$ for all $v \in V$), and
- cutting down the required separation to 1 (setting $d(vw) := 1$ for all $vw \in E^d$).

Clearly, this problem is a relaxation of the original frequency planning problem and an optimal solution of this problem yields a lower bound. All three simplifications together lead to a problem which is essentially the same as finding a minimum k-partition of the vertex set of the graph G, where edge-weights are derived from $d(vw)$ and c^{co}_{vw}.

A partition $V_1, \dots, V_k$ of V into at most $k = |F|$ many disjoint sets has to be determined such that the sum over the edge-weights in the induced subgraphs $G[V_i]$, $1 \leq i \leq k$, is minimized. The edge-weights are defined as follows. Set $M = \sum_{vw \in E^{co}} c^{co}_{vw} + 1$. The edge-weight function $\mu\colon E \mapsto \mathbb{R}_+$ is defined by $\mu_{vw} = c^{co}_{vw}$ in case $vw \in E^{co}$ and by $\mu_{vw} = M$ in case $vw \in E^d$. (Note that $E = E^{co} \cup E^d$ and $E^{co} \cap E^d = \emptyset$.)

There exists a feasible solution to the relaxed frequency planning problem if and only if a minimum k-partition has objective value less than M. Moreover, if the objective is less than M, its value is a lower bound for the original problem.

The min k-partition problem can be formulated as an integer linear program with an exponential number of constraints (cf. [20, Section 6.2]). An alternative formulation exploits the following properties of vectors of Euclidean length one (unit vectors) in $\mathbb{R}^n$, $n \geq 1$. First, for any k, $2 \leq k \leq n+1$, unit vectors $u_1, \ldots, u_k \in \mathbb{R}^n$ satisfying $\langle u_i, u_j \rangle \leq \delta$ for all $i \neq j$, it holds that $\delta \geq \frac{-1}{k-1}$. Second, there exist k unit vectors satisfying $\langle u_i, u_j \rangle = \frac{-1}{k-1}$ [24, 37]. Consider, for example, a k-simplex in $\mathbb{R}^n$, centered at the origin and scaled such that all vertices have a Euclidean distance of 1 to the origin. The vectors pointing at the vertices of that simplex have the desired property.

Hence, we may fix a set $U = \{u_1, \ldots, u_k\} \subset \mathbb{R}^{|V|}$ of unit vectors with $\langle u_i, u_j \rangle = \frac{-1}{k-1}$ for $i \neq j$ (and $\langle u_i, u_i \rangle = 1$ for $i = 1, \ldots, k$). These k vectors are used as labels (or representations) for the k partite sets in G. The min k-partition problem can then be phrased as follows. We search for a function $\phi \colon V \mapsto U$ that minimizes

$$\sum_{v,w \in V} \mu_{vw} \, \frac{(k-1)\, \langle \phi(v), \phi(w) \rangle + 1}{k} \tag{14}$$

Notice that the quotient in the summands evaluates either to 1 or to 0, depending on whether the same vector or distinct vectors are assigned to the respective two vertices.

If we assemble the scalar products $\langle \phi(v), \phi(w) \rangle$ into a square matrix X, being indexed row- and column-wise by V, then the matrix X has the following properties: all entries on the principal diagonal are ones, all off-diagonal elements are either $\frac{-1}{k-1}$ or 1, and X is positive semidefinite.

Notably, every matrix X satisfying the above properties defines a k-partition of V in the same way as ϕ does. This can be seen as follows. Since X is a positive semidefinite matrix, there exists a matrix C such that $X = C^T C$. We claim that C contains at most k distinct column vectors. For the sake of a contradiction, let us assume that $c_1, \ldots, c_{k+1}$ are $k+1$ distinct column vectors from C. Then $\langle c_i, c_j \rangle = \frac{-1}{k-1}$ for all $i \neq j$ and $\langle c_i, c_i \rangle = 1$ for all i has to hold. According to the above statement, however, the $k+1$ unit vectors may only have a mutual scalar product as low as $\frac{-1}{(k+1)-1} = \frac{-1}{k} > \frac{-1}{k-1}$. A contradiction. Therefore, the columns of C may indeed serve as the vectors assigned by ϕ, representing the partite sets.

The combinatorial problem to minimize (14) may be relaxed to a semidefinite program. First, the explicit reference to the set U is dropped, and the problem is rewritten as

$$\begin{aligned}
\min \quad & \sum_{v,w \in V} \mu_{vw} \, \frac{(k-1)\, X_{vw} + 1}{k} && && (15)\\
\text{s.t.} \quad & X_{vv} = 1 && \forall v \in V && (16)\\
& X_{vw} \in \{\tfrac{-1}{k-1}, 1\} && \forall v, w \in V && (17)\\
& X \succeq 0 && && (18)
\end{aligned}$$

Then, we replace the constraints (17) by $X_{vw} \geq \frac{-1}{k-1}$ (note that $X_{vw} \leq 1$ is enforced implicitly by X positive semidefinite and $X_{vv} = 1$). The SDP relaxation of the problem now reads

$$
\begin{aligned}
\min \quad & \sum_{v,w \in V} \mu_{vw} \frac{(k-1)\, X_{vw} + 1}{k} && && (19)\\
\text{s.t.} \quad & X_{vv} = 1 && \forall v \in V && (20)\\
& X_{vw} \geq \tfrac{-1}{k-1} && \forall v, w \in V && (21)\\
& X \succeq 0 && && (22)
\end{aligned}
$$

In [20], this bound is tested successfully for the COST 259 benchmark instances using recently developed software for semidefinite programming by Sturm [52], Burer, Monteiro and Zhang [14], and Helmberg [29].

7 Conclusions

Frequency assignment is an important problem in wireless communication with direct consequences for the quality of service and indirect economic impact, e. g., via the choice and utilization of equipment. There is not just one frequency assignment problem. The questions to be addressed depend strongly on the technology employed, government regulations, etc. Thus, modeling issues arise, and the mathematical problems to be solved vary from technology to technology, frequency band to frequency band, or even operator to operator. Focusing on the GSM technique, we have surveyed in this paper the mathematical questions arising in this context, many related to coloring problems in graphs. Practical frequency assignment problems are large and hard—not only in theory. Data from practice are available, so that the structure of practical instances can be studied. Although the utilization of optimization methodology has resulted in considerable progress compared with traditional planning techniques, the solution machinery is not in a satisfactory state yet. There are often large gaps between upper and lower bounds on the interference, for instance. Much better quality guarantees for the solution methods applied to instances from practice need to be achieved. The only path to progress is further research addressing the issues involved. That is, investigating the mathematical models described in this paper instead of continuing to study the traditional problems in coloring theory, a nontrivial challenge.

References

[1] K. I. Aardal, A. Hipolito, C. P. M. van Hoesel, and B. Jansen. A branch-and-cut algorithm for the frequency assignment problem. Research Memorandum 96/011, Maastricht University, 1996. Available at http://www.unimaas.nl/~svhoesel/.

[2] K. I. Aardal, C. P. M. van Hoesel, A. M. C. A. Koster, C. Mannino, and A. Sassano. Models and solution techniques for the frequency assignment problem. In preparation, Maastricht University, 2000.

[3] K. I. Aardal, C. A. J. Hurkens, J. K. Lenstra, and S. R. Tiourine. Algorithms for frequency assignment problems. Technical report, Technische Universiteit Eindhoven, 1999.

[4] S. M. Allen, D. H. Smith, and S. Hurley. Lower bounding techniques for frequency assignment. *Discrete Mathematics*, 197/198:41–52, 1999.

[5] L. G. Anderson. A simulation study of some dynamic channel assignment algorithms in a high capacity mobile telecommunications system. *IEEE Transactions on Communications*, 21:1294–1301, 1973.

[6] D. Applegate, R. Bixby, V. Chvátal, and W. Cook. On the solution of traveling salesman problems. In *Proceedings of the International Congress of Mathematicians Berlin 1998*, volume III of *Documenta Mathematica*. DMV, 1998.

[7] G. Ausiello, P. Crescenzi, G. Gambosi, V. Kann, A. Marchetti-Spaccamela, and M. Protasi. *Complexity and Approximation: combinatorial optimization problems and their approximability properties.* Springer-Verlag, 1999.

[8] M. Bellare, O. Goldreich, and M. Sudan. Free bits, pcps and non-approximability - towards tight results. *SIAM Journal on Computing*, 27:804–915, 1998.

[9] H. P. van Benthem. GRAPH generating radio link frequency assignment problems heuristically. Master's thesis, Delft University of Technology, 1995.

[10] M. Biró, M. Hujter, and Zs. Tuza. Precoloring extensions. I. interval graphs. *Discrete Mathematics*, 100:267–279, 1992.

[11] R. Borndörfer, A. Eisenblätter, M. Grötschel, and A. Martin. Frequency assignment in cellular phone networks. *Annals of Operations Research*, 76:73–93, 1998.

[12] F. Box. A heuristic technique for assigning frequencies to mobile radio nets. *IEEE Transactions on Vehicular Technology*, 27:57–74, 1978.

[13] D. Brélaz. New methods to color the vertices of a graph. *Communications of the ACM*, 22:251–256, 1979.

[14] S. Burer, R. D. C. Monteiro, and Y. Zhang. Interior-point algorithms for semidefinite programming based on a nonlinear programming formulation. Technical Report TR 99–27, Department of Computational and Applied Mathematics, Rice Unviversity, Houston, Texas 77005, dec 1999.

[15] A. Caminada. CNET France Telecom frequency assignment benchmark. URL: http://www.cs.cf.ac.uk/User/Steve.Hurley/f_bench.htm, 2000.

[16] K.-N. Chang and S. Kim. Channel allocation in cellular radio networks. *Computers and Operations Research*, 24:849–860, 1997.

[17] D. Costa. On the use of some known methods for t-colourings of graphs. *Annals of Operations Research*, 41:343–358, 1993.

[18] M. B. Cozzens and F. S. Roberts. T-colorings of graphs and the channel assignment problem. *Congressus Numerantium*, 35:191–208, 1982.

[19] G. Dueck and T. Scheuer. Threshold accepting: a general purpose optimization algorithm appearing superior to simulated annealing. *Journal of Computational Physics*, pages 161–175, 1990.

[20] A. Eisenblätter. *Frequency Assignment in GSM Networks: Models, Heuristics, and Lower Bounds.* PhD thesis, Technische Universität Berlin, Berlin, Germany, 2000.

[21] A. Eisenblätter and A. M. C. A. Koster. FAP web - A website devoted to frequency assignment. URL: http://fap.zib.de, 2000.

[22] P. Erdős, A. L. Rubin, and H. Taylor. Choosability in graphs. *Congressus Numerantium*, 26:125–157, 1979.

[23] M. Fischetti, C. Lepschy, G. Minerva, G. Romanin-Jacur, and E. Toto. Frequency assignment in mobile radio systems using branch-and-cut techniques. *European Journal of Operational Research*, 123:241–255, 2000.

[24] A. Frieze and M. Jerrum. Improved Approximation Algorithms for MAX k-CUT and MAX BISECTION. *Algorithmica*, 18:67–81, 1997.

[25] M. Grötschel, L. Lovász, and A. Schrijver. Polynomial algorithms for perfect graphs. *Annals of Discrete Mathematics*, 21:325–356, 1984.

[26] M. Grötschel, L. Lovász, and A. Schrijver. *Geometric Algorithms and Combinatorial Optimization.* Number 2 in Algorithms and Combinatorics. Springer-Verlag, 1988.

[27] W. K. Hale. Frequency assignment: Theory and applications. *Proceedings of the IEEE*, 68:1497–1514, 1980.

[28] M. Hellebrandt and H. Heller. A new heuristic method for frequency assignment. Technical Report TD(00) 003, COST 259, Valencia, Spain, Jan. 2000.

[29] C. Helmberg. SBmethod — A C++ implementation of the spectral bundle method, manuel to version 1.1. Technical Report 00-35, Konrad-Zuse-Zentrum für Informationstechnik Berlin (ZIB), Berlin, Germany, 2000.

[30] C. Helmberg. Semidefinite programming for combinatorial optimization. Technical Report 00-34, Konrad-Zuse-Zentrum für Informationstechnik Berlin (ZIB), Berlin, Germany, 2000. Habilitation TU Berlin.

[31] S. Hurley, D. H. Smith, and S. U. Thiel. FASoft: A system for discrete channel frequency assignment. *Radio Science*, 32:1921–1939, 1997.

[32] J. Janssen and K. Kilakos. Polyhedral analysis of channel assignment problems: (i) tours. Technical Report CDAM-96-17, London School of Economics, 1996.

[33] J. Janssen and K. Kilakos. An optimal solution to the "Philadelphia" channel assignment problem. *IEEE Transactions on Vehicular Technology*, 48(3):1012–1014, 1999.

[34] B. Jaumard, O. Marcotte, and C. Meyer. Mathematical models and exact methods for channel assignment in cellular networks. In B. Sansáo and P. Soriano, editors, *Telecommunications Network Planning*, chapter 13, pages 239–255. Kluwer Academic Publishers, Boston, 1999.

[35] D. S. Johnson. Worst case behaviour of graph colouring algorithms. *Congressus Numerantium*, 10:513–527, 1974.

[36] M. Jünger, G. Reinelt, and G. Rinaldi. *Network Models*, volume 7 of *Handbooks in Operations Research and Management Science*, chapter The travelling salesman problem, pages 225–330. Elsevier Science B.V., 1995.

[37] David Karger, Rajeev Motwani, and Madhu Sudan. Approximate graph coloring by semidefinite programming. *Journal of the ACM*, 45(2):246–265, 1998.

[38] M. Karoński. Random graphs. In R. Graham, M. Grötschel, and L. Lovász, editors, *Handbook of Combinatorics*, volume 1, chapter 6, pages 351–380. Elsevier Science B. V., 1995.

[39] R. M. Karp. Reducibility among combinatorial problems. In R. E. Miller and J. W. Thatcher, editors, *Complexity of Computer Computations*, pages 85–103. Plenum Press, New York, 1972.

[40] A. W. J. Kolen. A genetic algorithm for frequency assignment. Technical report, Maastricht University, 1999. Available at http://www.unimaas.nl/˜akolen/.

[41] A. M. C. A. Koster. *Frequency Assignment - Models and Algorithms.* PhD thesis, Maastricht University, 1999. Available at http://www.zib.de/koster/thesis.html.

[42] A. M. C. A. Koster, C. P. M. van Hoesel, and A. W. J. Kolen. Lower bounds for minimum interference frequency assignment problems. Technical Report RM 99/026, Maastricht University, 1999. Available at http://www.zib.de/koster/.

[43] R. Mathar and J. Mattfeldt. Channel assignment in cellular radio networks. *IEEE Transactions on Vehicular Technology*, 42:647–656, 1993.

[44] B. H. Metzger. Spectrum management technique, Fall 1970. Presentation at 38th National ORSA meeting (Detroit, MI).

[45] R. A. Murphey, P. M. Pardalos, and M. G. C. Resende. Frequency assignment problems. In D.-Z. Du and P. M. Pardalos, editors, *Handbook of combinatorial optimization*, volume Supplement Volume A. Kluwer Academic Publishers, 1999.

[46] A. Raychaudhuri. *Intersection Assignments, T-Colourings and Powers of Graphs.* PhD thesis, Rutgers University, 1985.

[47] ROADEF challenge 2001. URL: http://www.prism.uvsq.fr/˜vdc/ROADEF/CHALLENGES/2001/, 2000.

[48] F. S. Roberts. On the mobile radio frequency assignment problem and the traffic light phasing problem. *Annals of New York Academy of Sciences*, 319:466–483, 1979.

[49] F. S. Roberts. t-colorings of graphs: Recent results and open problems. *Discrete Mathematics*, 93:229–245, 1991.

[50] K.N. Sivarajan, R.J. McEliece, and J.W. Ketchum. Channel assignment in cellular radio. In *Proceedings of the 39th IEEE Vehicular Technology Conference*, pages 846–850, 1989.

[51] S. Stahl. n-tuple colorings and associated graphs. *Journal of Combinatorial Theory*, 20:185–203, 1976.

[52] J. F. Sturm. Using SeDuMi 1.02, a MATLAB toolbox for optimization over symmetric cones. Technical report, Communications Research Laboratory, McMaster University, Hamilton, Canada, 1998. Available at: http://www.unimaas.nl/˜sturm/.

[53] B. A. Tesman. *T-colorings, list T-colorings, and set T-colorings of graphs.* PhD thesis, Department of Mathematics, Rutgers University, 1989.

[54] B. A. Tesman. Set T-colorings. *Congrussus Numerantium*, 77:229–242, 1990.

[55] B. A. Tesman. List T-colorings. *Discrete Appied Mathematics*, 45:277–289, 1993.

[56] B. Toft. Colouring, stable sets and perfect graphs. In R. Graham, M. Grötschel, and L. Lovász, editors, *Handbook of Combinatorics*, volume 1, chapter 4, pages 233–288. Elsevier Science B. V., 1995.

[57] C. Valenzuela, S. Hurley, and D. H. Smith. A permutation based genetic algorithm for minimum span frequency assignment. *Lecture Notes in Computer Science*, 1498:907–916, 1998.

[58] V. G. Vizing. Critical graphs with given chromatic class (russian). *Diskret. Analiz*, 5:9–17, 1965.

[59] W. Wang and C. K. Rushforth. An adaptive local-search algorithm for the channel-assignment problem (CAP). *IEEE Transactions on Vehicular Technology*, 45:459–466, 1996.

[60] J. A. Zoellner and C. L. Beall. A breakthrough in spectrum conserving frequency assignment technology. *IEEE Transactions on Electromagnetic Compatiblity*, 19:313–319, 1977.