Konrad-Zuse-Zentrum
für Informationstechnik Berlin

Takustraße 7
D-14195 Berlin-Dahlem
Germany

MARTIN GRÖTSCHEL

Frequency Assignment in Mobile Phone Systems

ZIB-Report 00-58 (Dezember 2000)

Frequency Assignment in Mobile Phone Systems

Martin Grötschel

11. Januar 2001

Wireless Communication and Frequencies

Wireless communication technology is the basis of radio and television broadcasting, it is used in satellite-based cellular telephone systems, in point-to-multipoint radio access systems, and in terrestrial mobile cellular networks, to mention a few such systems (see [5] for more detailed information).

Wireless communication networks employ radio frequencies to establish communication links. The available radio spectrum is very limited. To meet today's radio communication demand, this resource has to be administered and reused carefully in order to control mutual interference. The reuse can be organized via separation in space, time, or frequency, for example. The problem, therefore, arises to distribute frequencies to links in a "reasonable manner". This is the basic form of the frequency assignment problem. What "reasonable" means, how to quantify this measure of quality, which technical side constraints to consider cannot be answered in general. The exact specification of this task and its mathematical model depend heavily on the particular application considered. In this paper we discuss this issue with respect to the GSM standard for mobile communication.

Mobile Cellular Networks, GSM

I will concentrate here 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. I will 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 serving several hundred million customers. The new worldwide standard UMTS (Universal Mobile Telecommunication System) is expected to become commercially available around 2002. It 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 GMS: frequency or time division are used in combination with code division multiple access (CDMA) technology.

Channel Spectrum

The typical situation in GSM frequency planning is as follows. A telephone company (let us call it the *operator*) has bought 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 usually denoted by $1, 2, \ldots, N$, 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 buys 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.

BTSs, TRXs and Cells

To serve his customers an operator has to solve a number of nontrivial problems. In an intitial 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. Each 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 this antenna. The first TRX handles the signalling 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 signalling 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.

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. A BSC is, e.g., in charge of the management of hand-overs.

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 intend to 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.

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 where the mobile phone is currently located. This is done via so called control channels. Whenever a communication demand arises, the system decides which TRX is going to handle the communication. This decision is based on the signal strengths of the various TRXs that are able to communicate with the phone as well as on the current traffic. The mobile phone is tuned to the channel of the TRX that presently appears to serve the phone best. 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 are not possible.

Interference

Whenever two cells overlap and use the same channel, interference (signal-to-noise ratio at the receiving end of a connection) 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 strongly influences 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, 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. They combine this with practical experience and extensive measurements. The result is an interference prediction model with which the so called *co-channel interference* that 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 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 behaviour? 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).

Separation and Blocked Channels

There are also hard constraints. 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 $i+1, i-1$. 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.

The situation is even more complex. Due to government regulations, agreements with operators in neighbouring 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*.

Interference Graph

A feasible assignment of channels to TRXs clearly has to satisfy all separation constraints. Blocked channels 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 $c^{ad}(vw)$.

Two "Natural" Approaches

A first attempt to solve the frequency assignment problem is obvious. We try to find a spectrum, i.e., a number of channels $1, \ldots, N$ such that the N available channels can be assigned to TRXs so that no interference occurs.

No interference is, of course, a good aim, but this task is unrealistic for several reasons. A mobile phone network is a "living system", i.e., new BTSs, antennas, etc., are installed regularly, old ones are replaced by new ones with different characteristics. It is impossible to change the spectrum each time the network changes. Moreover, the number of channels may be fixed or channels may only be available in bundles (i.e., one may buy 75, 100 or 125, but nothing else). Frequencies are expensive and cost reasons may require that some interference is tolerated. In fact, some interference may be unavoidable. We have data of mobile phone systems where the largest clique in the interference graph is about twice as large as the number of available channels and where the largest degree of a node is ten times as large as the number of available channels.

Another classic choice for the solution of the frequency assignment problem is the following. We choose a threshold value t and consider the graph $G^t = (V, E^t)$ where $E^t := \{ij \in E \mid c^{co}_{vw} \geq t\}$. Now we try to find the coloring number of G^t, or try to color G^t with N colors. In other words, we consider interference below t tolerable and try to find an assignment of channels to TRXs such that as few channels as possible are used (a color represents a channel, no two nodes with the same color are allowed to be adjacent) or we try to use the available channels so that no "high interference" occurs. Of course, if for a given threshold t no feasible coloring can be found, one has to modify t and try again.

This approach is unable to handle separation constraints and ignores adjacent channel interference. It was the "standard approach" in the early days of the mobile phone era but did not prove efficient in the more complex environment we have today.

Minimizing Interference

There are several other ways of modelling the frequency assignment problem mathematically, see [5]. For reasons of brevity I will focus on an approach that was employed in a joint project of the Konrad-Zuse-Zentrum and E-PLUS, one of the four operators in Germany, and which has resulted in very satisfactory channel assignments.

Let $G = (V, E)$ be the interference graph introduced before. Let $C = [1, \dots, N]$ be the set of available channels, and let, for each TRX $v \in V$, B_v denote the subset of channels blocked at node v. The values $c^{co}(vw)$, $c^{ad}(vw)$ denote, for each edge $vw \in E$, the co-channel and adjacent-channel interference arising when TRXs v and w operate on the same or on adjacent channels. Moreover, let $d(vw) \in \mathbb{Z}_+$ denote the separation necessary between the channels assigned to TRXs v and w. Thus, the input to a frequency assignment problem is a 7-tuple $(V, E, C, \{B_v\}_{v \in V}, d, c^{co}, c^{ad})$, briefly called *network* here. A *frequency assignment* for the network is a function $y : V \to C$. It is called *feasible* if it satisfies the following side constraints

$$y(v) \in C \backslash B_v \text{ for all } v \in V$$

$$|y(v) - y(w)| \geq d(vw) \text{ for all } vw \in E$$

The objective is to minimize the sum of co- and adjacent-channel interference, more formally:

$$\min_{y \text{ feasible}} \sum_{\substack{vw \in E: \\ y(v)=y(w)}} c^{co}(vw) \quad + \sum_{\substack{vw \in E: \\ |y(v)-y(w)|=1}} c^{ad}(vw) \tag{FAP}$$

This version of the frequency assignment problem is a generalization of list colorings in graph theory and it is related to the well known T-coloring problem.

There are several ways to reformulate (FAP) in terms of other standard models of combinatorial optimization, e.g., there is a stable set model and a so called orientation model which is related to linear ordering.

Several modifications of FAP have to be considered in practice. No operator wants to change all channel assignments whenever a new plan has to be computed. Some assignments have to stay fix (that is easy to achieve), sometimes one looks for the smallest number of channel adjustments within a certain range of interference, or one requires that, e.g., at most 100 of the assignments are changed.

FAP is difficult in terms of complexity theory. Deciding whether a TRX network allows a feasible assignment is $\mathcal{NP}$-complete; the optimization problem is strongly $\mathcal{NP}$-hard.

FAP is also difficult in practice. Nobody can solve realistic instances to optimality. Satisfactory lower bounds on the objective function value are very hard to obtain. All approaches based on polyhedral combinatorics and linear programming have failed so far. There is some hope to exploit semidefinite relaxations of FAP.

The whole available "zoo" of heuristics has been tried for the solution of FAPs. Considerable improvements over previous approaches can be achieved. There are some spectacular successes, but at present, the gap between lower and upper bounds – computable in practice – is still very large.

In my talk on this subject I will elaborate on the mathematical modelling of the FAP, on the development of heuristics and on the approaches with which lower bounds have been computed. I will present examples from practice that show what can be achieved today and how this mathematical approach compares to more traditional planning techniques.

This lecture is based on joint work of the telecommunications group at the Konrad-Zuse-Zentrum, particular on the work of Andreas Eisenblätter [3]. Further references are [1], [2], [4]. The FAP website [5] is another excellent source of information.

References

[1] R. Borndörfer, A. Eisenblätter, M. Grötschel, and A. Martin, "Frequency Assignment in Cellular Phone Networks", *Annals of Operations Research*, 76:73-93 (1998).

[2] R. Borndörfer, A. Eisenblätter, M. Grötschel, and A. Martin, "The orientation model for Frequency Assignment Problems", Technical Report TR 98-013, Konrad-Zuse-Zentrum für Informationstechnik Berlin, (1998a).

[3] A. Eisenblätter, “Frequency Assignment in GSM Networks: Models, Heuristics, and Lower Bounds”, Ph.D. Thesis, TU Berlin 2000, to appear.

[4] A.M.C.A. Koster, “Frequency Assignment – Models and Algorithms”, Ph.D. Thesis, Universiteit Maastricht, Maastricht, The Netherlands (1999).

[5] FAP Web – A Website about Frequency Assignment Problems, `http://fap.zib.de`

veröffentlicht in: *Lecture Notes in Computer Science* 1974, Kapur, Prasat