Wiesbaden '95 revisited
– T-DAB planning parameters, reference networks and frequency-planning algorithms

T. O'Leary
EBU Technical Department

As a follow-up to the two previous articles in EBU Technical Review about Wiesbaden '95 [1][2], the author describes the EBU Systhesis computer program that was developed for allotting T-DAB frequency blocks during the Wiesbaden planning process.

1. Introduction

The CEPT [1] held a T-DAB [2] Planning Meeting in Wiesbaden during three weeks in July 1995 (referred to as Wi'95 hereafter). There was a great deal of preparatory work by both the CEPT and the EBU, which was carried out before the actual distribution of frequency allotments began.

To a large degree, the planning of T-DAB at Wiesbaden was different than at some of the other recent terrestrial planning conferences held in Region 1, e.g. for VHF/FM in 1984 or the African Television Conference in 1989. In particular, because of the newness of T-DAB and its digital characteristics, interference assessment and coverage planning were concepts which had to be restructured, if not reinvented, as the planning process progressed. Indeed, at the beginning of the planning studies, there were not even any frequency bands available for the proposed new service; it was only during the preparatory period that a gradual outline of suitable frequencies, the necessary number, and their (eventual) availability, became clear.

Because of the additional dearth of information concerning transmitter sites and related radiation characteristics, it was decided quite early to rely on allotment planning methods, rather than the usual individual-assignment planning methods. This need, plus the novel "properties" of wideband digital systems, gave rise to the development of the idea of a "reference network", which exploits the single frequency network (SFN) concept. This concept derives from the digital characteristics of T-DAB which allow identical signals from two (or more) transmitters to "add" in many circumstances, thus providing a stronger signal than either (or any) would provide alone. In this way, it was possible to plan for full wide-area coverage using a single frequency for any given allotment area.

In order to start a "viable" T-DAB service in any country, it was felt that the availability of at least twelve programmes would be imperative at all receiving sites. Thus, two distinct T-DAB

1. European Conference of Postal and Telecommunications Administrations.

2. Terrestrial Digital Audio Broadcasting.

Original language: English
Manuscript received: 15/12/98.

blocks (each providing six high quality programmes) were found necessary to provide the minimum number of programmes in all allotment areas in Europe. Because of the propensity for interference to arise between different T-DAB signals on the same frequency (in this respect, T-DAB is not unlike any other broadcasting service), it was necessary to have access to a sufficient number of T-DAB frequency blocks so that, through physical separation of the co-frequency allotments, the predictable disruption of service could be avoided.

As mentioned above, no T-DAB dedicated band was available for planning this new service. It was thus incumbent on the CEPT and frequency planners to "make room" for T-DAB in the already overfilled broadcasting bands. This initially lead to the selection of certain "preferred" T-DAB blocks in Bands I and III (television channels 2 to 12); eventually, due to extreme need, a new "channel" – the so-called "channel 13" (230 - 240 MHz) – was also permitted in some countries for T-DAB use. Furthermore, the 1.5 GHz band was partitioned in order to allow T-DAB services in the lower 15 MHz (1452 - 1467 MHz).

Needless to say, this pot-pourri of frequency possibilities, scattered over nearly 1500 MHz, lead to enormous planning difficulties – especially since, in all the bands considered, large numbers of "other" services were often present. Thus, suitable compatibility-sharing criteria had to be developed and, then, respected.

At Wi'95, the compatibility question was treated in two phases. The first phase, compatibility between T-DAB and other services, was solved by applying the agreed compatibility criteria and thus avoiding interference by reducing the number of T-DAB blocks that could be allotted to any given T-DAB requirement. The remaining "available" blocks were then treated during the second phase, whereby they were allotted on the basis of non-interference between T-DAB services. This was done with the help of a computerised allotment algorithm which will be described below.

Since Wi'95, there has been growing interest in providing a third allotment in certain CEPT countries. To this end, and to ensure a fair distribution of frequencies to all CEPT countries, the CEPT has called for another T-DAB Planning Meeting, to be held in the near future. For this new allotment, the CEPT has decided to use seven additional blocks in the 1.5 GHz band, situated adjacent to, and immediately above, those blocks at 1.5 GHz already used at Wi'95.

It may be thought that the next round of allotments will be quite easy in view of the preparations and planning already carried out in the years leading up to the Wiesbaden Planning Meeting in 1995. However, this article will point out a few of the remaining or new problems which have to be solved, as well as giving an overview of some of the computer planning algorithms used for the allotment process.

2. Planning parameters

Most of the planning parameters needed for the 1.5 GHz band were already established and agreed at Wi'95. However, in the aftermath of the conference it has been determined that the receiver noise figure should be about 3 dB higher than what was originally planned for. Certain adjustments during the T-DAB implementation phase, that is when establishing the actual transmitter networks, will be necessary to compensate for this increase. But to ensure a more efficient planning for future allotments using the additional T-DAB blocks (i.e. in the year 2000?), it will be necessary to correct this value.

The protection ratios for T-DAB versus T-DAB will probably remain as originally agreed but additional protection ratios will undoubtedly have to be determined vis-à-vis new "other services” which will be using the same (or nearby) frequencies, as soon as these other services and their technical characteristics become known.

3. Reference networks

Due to the differing propagation conditions between VHF and 1.5 GHz, a certain advantage may be seen in using the 1.5 GHz band for “smaller” coverage areas. In fact, for Wi’95, two separate reference networks were established for VHF and 1.5 GHz: a 60 km transmitter separation for VHF and a 15 km separation for the 1.5 GHz reference network. This 15 km separation distance for the 1.5 GHz band gives rise to about a 30 km coverage radius (see *Fig. 1*).

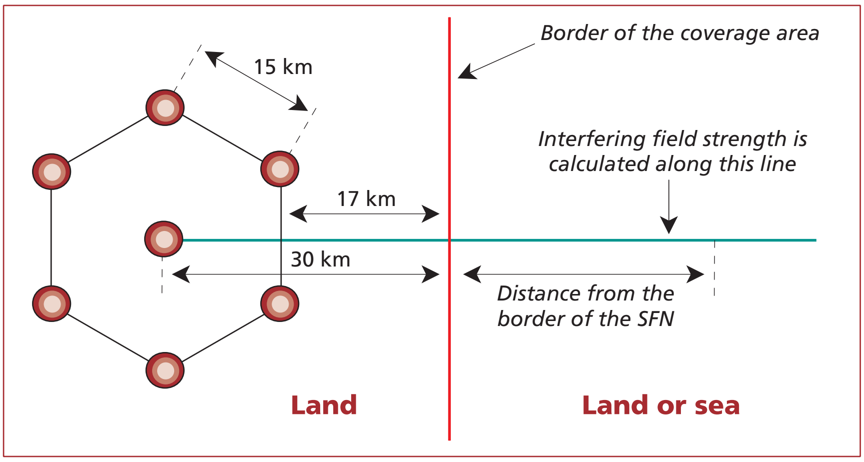

Figure 1
Wi’95 reference network for the 1.5 GHz band.

At present, it is being considered whether the reference network shown in *Fig. 1* should be modified, for at least two reasons. Firstly, as noted in *Section 2*, the original receiver noise figure used at Wi’95 has been found to be too low. A necessary increase of 3 dB would require some modification to the network in order to reflect the related increase in minimum equivalent field strength. In addition, due to a recent increased interest in smaller-area local services, it may be wise to perhaps foresee an additional choice (or choices) of reference network(s) in order to increase spectrum utilization in such cases. Much further work is needed in this area.

4. Frequency assignment algorithms

4.1. General

Frequency assignment is very simple, on the surface. All that must be done is to associate particular frequencies with particular service requirements, making sure that no two requirements having the same frequency are incompatible with each other. The problems arise when there are only a limited number of distinct frequencies to work with. It is similar to the task of covering the smallest area with a set of randomly-shaped tiles, using all of the tiles (only it’s harder). This might be relatively easy if all the tiles had the same regular shape and could thus be fitted together with no gaps, but it requires many attempts and manipulations if the tiles’ shapes are sufficiently irregular (think of a jigsaw puzzle where half of the pieces have been replaced by pieces from other puzzles!). It helps if you can find a set of rules (an “algorithm”) which has some common-sense basis.

Nowadays, problems needing lots of calculations or manipulations are usually solved with high-speed, high-capacity computers. And the frequency assignment problem, with its many intricate sets of compatibilities and incompatibilities, is one of them!

Much has been written about frequency assignment algorithms [3]. Methods range from complete exhaustive searches, to highly theoretical graph-theoretic methods (applying efficient mathematical shortcuts to examine subsets of a complete search), and from Monte Carlo methods [4] to intuitive methods. Even though these approaches (except the first) examine only a minuscule subset of the total possibilities, their application can often require large computation times.

Although the quest for rapid, efficient algorithms has been underway for many years (the new Holy Grail?), it seems that at least some frequency planners are not yet completely satisfied – because the quest still continues. The basic reason for this is the fact that many algorithms are "situation oriented", i.e. they are "optimal" for solving certain types of problems, but not for others. Even a highly-developed algorithm (often with relatively long computer runtimes) may find a solution inferior, on occasion, to that found by a less highly-developed algorithm which, at the same time, is also often quicker. It is basically for this reason that the EBU has developed a frequency-synthesis program which consists of a large number of relatively simple (and thus fast) algorithms which (hopefully) will yield a (near) optimal solution quickly [3].

4.2. *Intuitive description of an assignment process*

4.2.1. It all adds up

When children learn basic addition at school, they learn it as a sequential binary process (not digital binary, but rather operational binary). That is, they learn to add the first two numbers of the set to be added (the addends), take the result and add it to the next addend, and so on until the final sum is reached. It would be much too complicated in general to add all the addends (if there are more than 2) simultaneously. (For some kids it's even too difficult to do so with only two addends!) Likewise, the assignment procedure used in the EBU algorithms is a sequential approach, making one assignment at a time. Of course, just as the addition of many addends may be done more simply by adding selected pairs first (those summing to multiples of 10, for example), it may be propitious to associate (or not) particular requirements at the beginning, or thereafter, as appropriate. This requires a certain in-sight (or luck) to be able to know which are the simplifying associations.

4.2.2. Boxing stones

To give an analogy, which is really not that much different from the frequency assignment process, consider the task of filling a given box (or a set of boxes) with a large number of stones of different sizes and shapes. One method to do this would be to put random stones into random boxes and hope that no stones are left over at the end. (This approach could be carried out "all at once" and then could be called the "avalanche approach"!).

3. As a result of experience at Wi'95, it could actually be questioned whether a complicated, long running algorithm would have really been required at all during Wi'95. For the first two and a half weeks (out of a total of three), the number of requirements exceeded by far the capacity of the available T-DAB frequency blocks. In this case, even a complete search would not have yielded a solution.

4.2.3. Ordering stones

Another "common sense" approach might be to first put the large boulders into the boxes and to fill the remaining interstices with the smaller stones and pebbles. Once again, it would not be possible to put all the stones into the boxes simultaneously in the correct space-saving manner. Usually this would require a certain amount of pre-reflection, or even post-manipulation, at each step. Thus, a sequential approach suggests itself here also. To do this in an orderly common-sense fashion, the stones could first be sorted according to size (largest to smallest) and then boxed in that order. This could be called the "Largest First" approach. Doing the ordering in the opposite direction could be called the "Smallest First" approach. Even here, there might be reasons to "jump ahead" sometimes and take a nice subset of pebbles which all fit together and put them, as a whole, into a hole into which they just fit, precisely. Otherwise, if the pebbles wait their normal sequential turn, they might be dispensed in a less space-saving fashion.

4.2.4. Ordering boxes

A strategy might also be developed for determining the order in which the boxes are to be filled. For example, it might be efficient to fill the first box as far as possible before starting the second box, then filling the second box as far as possible before starting the third, and so on, only returning to previous boxes when the sizes of the stones have reached a point where they may be put into the remaining (small) gaps. This could be called the "First Available" approach. Another strategy might be to fill the boxes in a (relatively) uniform way so that, at any given step, a stone is placed in the box (having sufficient space, of course) in which the least number of stones has already been attributed. This could be called the "Least Heavily Occupied" approach. Yet another strategy might be to load, preferentially, those boxes already having the largest number of stones, when possible. This could be called the "Most Heavily Occupied" approach.

Still another more complicated, approach would be to determine at each step, as the assignment process progresses, how many of the remaining stones would fit individually into each box, assuming that none of the other remaining stones were to be taken into account. We could say that entry to a given box is "requested" by a certain number of stones. And then we could attribute stones to boxes according to a "Most Heavily Requested" or "Least Heavily Requested" principle. The reader may have noticed that although these two approaches are also simple, it requires a little bit of extra work after each assignment to re-evaluate the "occupancy" and "requestedness" of the boxes, i.e. there is a continual change which may or may not alter the future ordering of the boxes.

4.2.5. Chaos out of orders

The next simple extension would be to order the selection of the stones and also to order the selection of the boxes. In *Sections 4.2.3* and *4.2.4* above, two possibilities for ordering stones and five possibilities for ordering boxes, respectively, were described. If we take all possible combinations of stone order and box order, we arrive at 2 x 5 = 10 possible "filling algorithms".

Furthermore, we might order the stones not only on the basis of size but rather on the weight, or the colour or whatever of the stones, if it were thought that these properties were associated somehow with efficient packing. We then could take all possible combinations of stone and box orderings which could lead, some might say to chaos, but in any case to a large number of "filling algorithms", depending only on our inventiveness. Fortunately this "chaos" can be mastered by a computer which is especially suited for such intricacies.

It is possible to evaluate the results of resolving this "chaos". If the "filling" algorithms are efficient (or at least one of them is) there should be no stone which has not found an appropriate box (assuming of course that such a solution exists at all). Those algorithms which accommodate a larger number of stones are "better" than those which only accommodate fewer.

If there is no single algorithm which is "best" all the time, it is obvious that using more than one is preferable. An algorithm which is never good can be disregarded. To determine the "quality" of the algorithms it is necessary to carry out many tests on many configurations. Of course, if a configuration is very complicated, the optimum result may not be known (if it were, then the algorithms wouldn't be needed!) and so the extent to which an algorithm is "perfect" cannot always be precisely judged.

4.3. Outline of the general approach

To get back to frequency assignment for a moment, we can make the above analogy more meaningful if we associate "largeness of stone" with "incompatibility of requirement". Here we mean the global incompatibility of the requirement, i.e. the total number of incompatibilities that it has with all the other requirements. After all, it is this incompatibility which prevents the relevant requirements from sharing a given frequency. This association also indicates why frequency assignment is so much more complicated when compared with putting stones into boxes. "Largeness" in frequency assignment is related to the existence of the other requirements, and the "boxes" can become smaller with time. For example two "large" requirements may be able to fit into the same box (i.e. have the same frequency) , whereas one "small" one may not, because of previous frequency assignments.

To give a pictorial representation of the global incompatibility situation, use is often made of a graph consisting of vertices (the requirements) and lines (called edges) connecting some of those vertices. Two vertices are connected by edges if the corresponding requirements are incompatible. See *Fig. 2* where, for example, requirements "A" and "B" are incompatible (there is an edge between "A" and "B") whereas requirements "A" and "H" are compatible (no edge between "A" and "H").

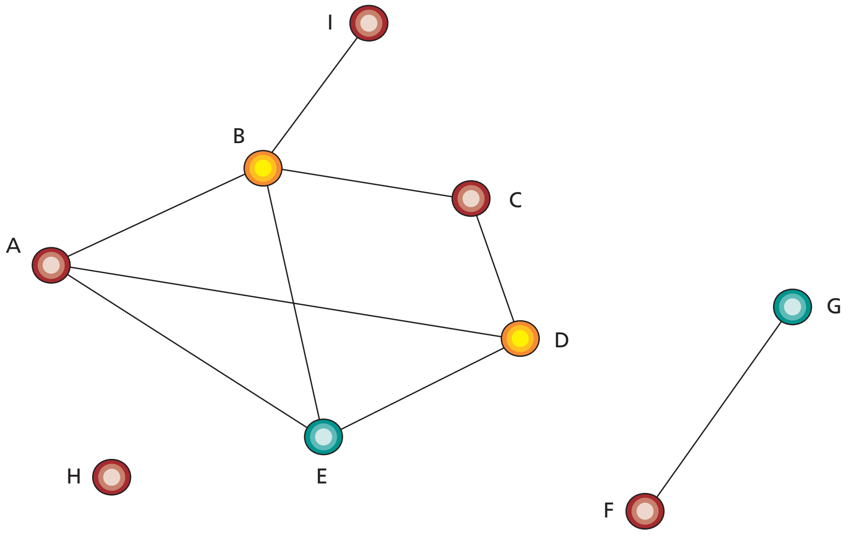

Figure 2
Incompatibility graph for 9 requirements.

It should be clear visually from *Fig. 2* that requirements "A", "D" and "E" cannot share the same frequency; indeed three distinct frequencies would be necessary just for these three requirements. A further inspection shows that three frequencies α, β and γ would be necessary to make a complete assignment as shown in *Table 1*.

Of course, many other 3-frequency distributions can also be found with a few seconds of thought. A more complex diagram with over 700 vertices (i.e. requirements)[4] would take a little longer!

This graphical type of representation is not suitable for computer manipulations, but can be made so by using an equivalent matrix representation as shown in *Table 2*. Here, incompatibilities between two requirements, X and Y, are indicated by an entry "1" at the intersection of the Xth row and the Yth column, and at the intersection of the Xth column and Yth row (for the purposes of ease, symmetry around the main diagonal is introduced). Similarly, compatibility is indicated by a "0" entry. A computer can easily be told how to manipulate sets of "0s" and "1s" such as these.

Requirement	Frequency
A	α
B	γ
C	α
D	γ
E	β
F	α
G	β
H	α
I	α

Table 1
Frequency assignment table for *Fig. 2.*

By summing the 0s and 1s in a given row, one can find the total number of incompatibilities associated with the corresponding requirement. For example, row "A" sums to 3, row "B" sums to 4, etc. These values indicate the "largeness" of the global incompatibility of the requirement, and can be used to effect assignment ordering, e.g. "B", "A", "D", "E", "C", "F", "G", "I", "H" in descending order ("Largest First"), or "H", "I", "G", "F", "C", "E", "D", "A", "B" in ascending order ("Smallest First"). Experience has shown that a "Largest First" approach contributes far more often to good assignment results than does a "Smallest First".

However there are also many other ways of ordering these requirements. For example, a "Largest First" and "Smallest Last" ordering would be equivalent if the counting is done as in *Table 2*. But a different type of "Smallest Last" can also be devised if, once the "smallest" requirement is determined and put at the end of the list, its contributions to the interference weight of the remaining requirements are deleted, and a new ordering of these requirements is determined on this basis. Then the new "smallest" is placed at the next-to-last position on the ordering list, and so on. This is called "Dynamic Smallest Last" ordering. Although, at first sight, this may seem a trivial modification, in fact it can lead to substantial reorderings which often contribute to "very good" assignment results. Similarly a "Dynamic Largest First" ordering can also be envisaged.

	A	B	C	D	E	F	G	H	I	SUM
A	0	1	0	1	1	0	0	0	0	3
B	1	0	1	0	1	0	0	0	1	4
C	0	1	0	1	0	0	0	0	0	2
D	1	0	1	0	1	0	0	0	0	3
E	1	1	0	1	0	0	0	0	0	3
F	0	0	0	0	0	0	1	0	0	1
G	0	0	0	0	0	1	0	0	0	1
H	0	0	0	0	0	0	0	0	0	0
I	0	1	0	0	0	0	0	0	0	1

Table 2
Matrix representation of the incompatibility graph shown in Fig. 2.

4. As was the situation at Wi'95.

In the same way, other requirement orderings can be devised, based on other considerations, such as including incompatibilities (due to prior assignments) with available channels, when performing the ordering evaluation.

With respect to requirement ordering, a comment should be made here. We first define a term often used in the more technical literature. A "clique" is a set of requirements each of which is incompatible with all others in the set. The "size" of the clique is the number of requirements belonging to it. *Fig. 3* gives a graphical representation of a clique of size 5.

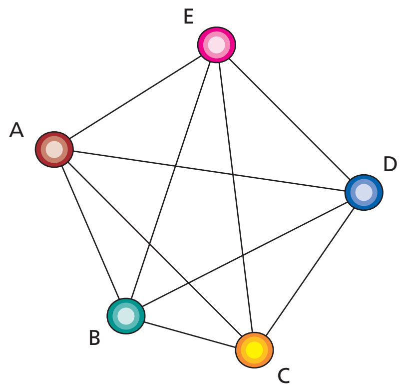

Figure 3
Clique size of 5.

It is clear that such a clique would require 5 distinct frequencies (and no fewer) in order to obtain an assignment plan with no incompatibilities. Thus, if we can determine the largest clique contained within a set of requirements, we know right away the minimum number of distinct frequency channels which will be needed. Knowing the members of the largest clique, it seems obvious that a channel assignment for these should be made at first [5], because otherwise other blockages may occur which could force the minimum number of needed channels to increase. This may happen anyway, as the example in *Fig. 4* shows. With a little reflection, it should be clear that two frequencies (corresponding to the maximum clique size) will not suffice [6].

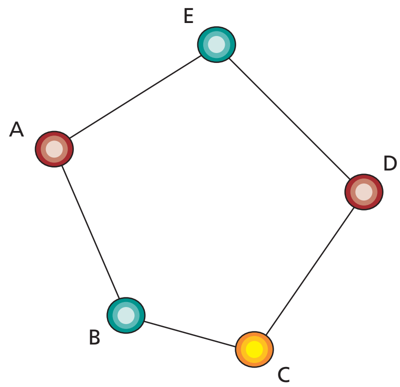

Figure 4
Clique size of 2:
frequencies needed = 3.

The methods of ordering the frequencies (called "boxes" in *Section 4.2*) also can be extended in similar ways. For example, not only could one use "First Available", "Least Heavily Occupied", "Most Heavily Occupied", "Least Heavily Requested" or "Most Heavily Requested", we can also use combinations such as "Least Heavily Occupied" and "Least Heavily Requested" etc.

One problem with a sequential approach, already alluded to in *Section 4.2*, is that what you do at one point in the assignment process will undoubtedly affect your possibilities at a later point. Sometimes this problem is addressed by using "backtracking" techniques [5]. The EBU algorithms also apply such techniques but only to a limited extent. The disadvantage is that, the more you backtrack (i.e. "undo", successively, previous assignments), the more time-consuming the algorithm will become, in general. And since a "bad" decision may have been taken at a relatively early point in the assignment process, effective backtracking could be very costly in terms of computer runtimes.

Another form of backtracking might be called "forward looking". "Forward looking", in its most obvious application, can be just as wasteful of computer runtimes as backtracking (they are almost the same procedures, really). However, other forms of "forward looking" can also be envisaged. The EBU algorithms also include such approaches. One of the methodologies uses a probabilistic formulation to estimate the probable "amount of damage" that may be

5. In fact, the best assignment results usually arise when making these allotments first.

6. This is easily seen as follows. If frequency α is assigned to "A", it cannot also be assigned to either "B" or "E". Assigning the second frequency β to "B" and "E", we find that "C" and "D" – being incompatible with B and D respectively – must be assigned frequency α; but "C" and "D" are mutually incompatible and thus cannot have the same frequency. Therefore a third frequency γ must be used.

done to future possibilities, when making one assignment choice as compared to another choice. Still another methodology treats the processed assignment channelling, and the list of waiting assignments, as a type of “electrical circuit” whereby a “least resistance” calculation is made and assignments are performed on this basis. These and other methodologies which have been devised are of course somewhat more demanding of computer time but, nevertheless, are often quite effective.

Further work is also being devoted to developing still other, perhaps more efficient, allotment algorithms.

4.4. Complications

Taking all possible combinations of the ordering procedures described in the previous paragraph leads to around 1000 distinct algorithms in the EBU Synthesis program. The running time for about 750 requirements at Wi’95 was about three hours, using a PC with a 66 MHz speed.

One complication, that will only be mentioned here, is that of the possibility of adjacent channel interference. The adjacent channel protection ratio is between – 30 and – 35 dB for both the VHF and 1.5 GHz bands. This means that no problems would arise when transmitting adjacent channels from the same site (not necessarily on the same antenna, though), as long as the powers were close enough together in magnitude. Problems could arise in overlapping service areas served by adjacent channels, at points where the difference in power levels becomes sufficiently large. Possible problems of this nature can be taken into account during the conversion process, when implementing T-DAB stations within an allotment area. If the adjacent channel protection ratios were sufficiently large, however, there would be a need to take this into account during the frequency allotment process itself. This could be done, if necessary, by making slight modifications to the algorithms described in the preceding section.

However, further unexpected complications can also arise (and did arise during Wi’95). For example, because channel 13, having only 10 MHz spectrum available, was not quite large enough to cater for six T-DAB blocks in the usual way, it was decided to reduce the guard band between the blocks labelled “13C” and “13D” in order to make room. This meant that two allotment areas “near” one another could not be assigned block 13C and block 13D, respectively. Such an additional restriction led to the need to modify the algorithms (under extreme pressure, as this restriction arose and was decided upon only during the Planning Meeting itself). It is hoped that this sort of “short notice complication” will not occur during the next session.

__Terry O’Leary__ received a doctorate in Physics at the University of California. In 1975, he joined the Institut für Rundfunktechnik (IRT Munich) where he conducted research on a range of topics including propagation, antennas, and terrestrial network and satellite planning.

In 1979, Dr O’Leary joined the EBU Technical Department where he became involved in many projects within the framework of EBU Working Party R. From 1984 to 1990, the IFRB benefitted from his specialist knowledge of HF and television network planning. He returned to the EBU in 1990 and was involved in T-DAB planning, WARC’77 BSS replanning and other projects.

Terry O’Leary is currently working with EBU Project Group BI TAPI (Terrestrial Allotment Plan Implementation), preparing for the next CEPT DAB Planning Meeting.

5. Conclusions

Although most of the necessary investigations that were completed in the time leading up to Wi'95 will still be valid, as can be seen from the above discussion, there is still much further work to be done in preparation for the next, 3rd Priority, T-DAB allotment plan. Although the time is short (less than two years?) there should still be enough to be well prepared, if this work begins as soon as possible.

Bibliography

[1] K.J. Hunt: **Personal reflections on the CEPT T-DAB Planning Conference: Wiesbaden, 3 – 21 July 1995**
EBU Technical Review No. 265, Autumn 1995.

[2] K.J. Hunt, T. Cesky et al.: **The CEPT T-DAB Planning Meeting – Wiesbaden, July 1995**
EBU Technical Review No. 267, Spring 1996.

[3] A. Ganst: **Computational complexity of some interference graph calculations**
IEEE Transactions on Vehicular Technology, Vol. 1, N.2, May 1990.

[4] T. Cesky: **Computerised VHF-FM frequency assignments using a Monte-Carlo method**
EBU Technical Review No. 231, October 1988.

[5] S.P. Leung: **Application of the Partial Backtracking Technique to the Frequency Assignment Problem**
IEEE Electromagnetic Compatibility Symposium, August 1981.