OFDM receivers

— impact on coverage of inter-symbol interference and FFT window positioning

Roland Brugger
IRT

David Hemingway
BBC

This article offers a general overview of the possible strategies for FFT window synchronization in OFDM receivers. These strategies are equally applicable to the T-DAB and DVB-T broadcasting systems.

The digital broadcasting systems DVB-T and T-DAB commonly exploit the single-frequency network (SFN) technique. In such networks, the treatment of inter-symbol interference and the synchronization strategy of the receivers is a crucial aspect in planning the networks and achieving the desired coverage. This is valid in two senses.

- Firstly, the performance of the receiver is strongly dependent on the way in which it positions the FFT window relative to the several received signals that can be present in a multipath environment or in an SFN. The position of the FFT window affects the receiver's behaviour with regard to inter-symbol interference (ISI).
- Secondly, the modelling of receiver behaviour in network coverage simulations has to be in line with the synchronization strategies and the treatment of inter-symbol interference in the receivers, in order to give reliable coverage predictions. However, the synchronization strategies of individual manufacturers are commercially sensitive and not publicly available. Therefore, predictions have to be carried out on the basis of assumptions.

1. Inter-symbol interference

1.1. General

In OFDM, the information is carried via a large number of individual carriers in a frequency multiplex. Each carrier transports only a relatively small amount of information, and high data capacities are achieved by using a large number of carriers within a frequency multiplex. The individual carriers are modulated by means of phase-shift and amplitude-modulation techniques. Each carrier has a fixed phase and amplitude for a certain time duration, during which a small portion of the information is carried. This unit of data is called a **symbol**; the time it lasts is called the **symbol duration**. After that time period, the modulation is changed and the next symbol carries the next portion of information. Modulation and demodulation are accomplished by the use of Inverse Fast Fourier Transformation (IFFT) and Fast Fourier Transformation (FFT) respectively. The symbol duration time is the inverse of the carrier spacing, ensuring orthogonality between the carriers.

In general, signals arriving at a receiver by different paths show different time delays which result in **inter-symbol interference** (ISI), a degradation in reception. An OFDM system with a multipath capability allows for the constructive combination of these signals. This is achieved by inserting a **guard interval** – a cyclic prolongation of the useful symbol duration of the signal. The FFT window, i.e. the time period for the OFDM demodulation, is then positioned in such a way that a minimum of inter-symbol interference occurs. This mechanism – as far as it is of interest for coverage predictions and network planning – is described in more detail in the following paragraphs.

In order to demodulate the signal – and looking at only one carrier – the receiver has to evaluate the symbol during the symbol duration. Three consecutive symbols in time, denoted by n–1, n and n+1, and the setting of the FFT window such that symbol n is evaluated by the receiver, are shown in *Fig. 1*. No guard interval is used in this example, and the FFT window has the same duration as the symbol.

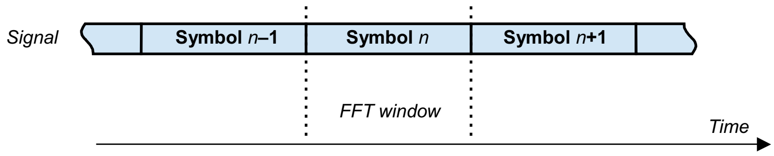

Figure 1
OFDM symbol duration and FFT window (no guard interval)

In an environment where several useful signals – either from multipath echoes or from other transmitters in an SFN – are available to the receiver, things become more complex. Usually, the signals arrive at different times at the receiver which, in the absence of a guard interval, makes correct synchronization to all the signals impossible. Such a situation, with two signals as an example, is depicted in *Fig. 2*. Synchronization to symbol n of signal 1 leads to an overlap of the FFT window with the preceding symbol n–1 of the delayed signal 2. Since this symbol n–1 carries different information from symbol n, the overlap acts as inter-symbol interference to the evaluation of symbol n.

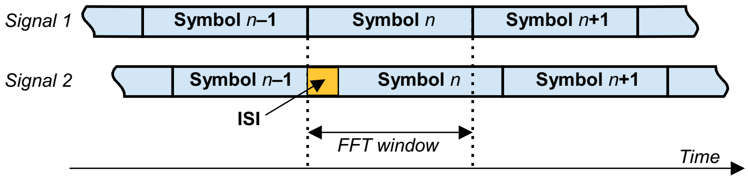

Figure 2
Inter-symbol interference with a delayed signal (no guard interval)

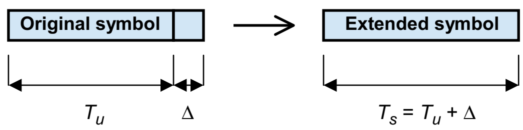

Figure 3
Increase of the symbol duration by the guard interval

In order to overcome the inter-symbol interference problem in DVB-T and T-DAB, part of the symbol is copied from the beginning of the symbol to the end, increasing its duration by a certain amount of time called the guard interval. This cyclic prolongation of the original symbol is shown in *Fig. 3*. The guard interval is denoted by Δ.

The new increased symbol duration is denoted by T_s and the original symbol duration is often called the **useful symbol duration** T_u. The duration of the FFT window during which the symbol is evaluated is kept at the original value T_u. The orthogonality relationship is kept with the original symbol duration T_u, not the extended T_s.

The improvement that is achieved by the insertion of the guard interval can be seen from *Fig. 4* with two signals as an example. The guard interval now allows for the FFT window to be positioned so that there is no overlap with a preceding or subsequent symbol, thus avoiding ISI.

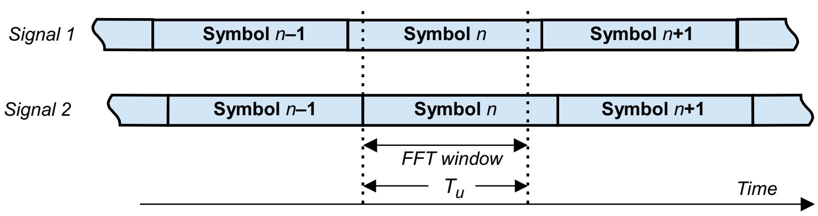

Figure 4
Guard interval utilisation

The fact that the duration of the FFT window is now smaller than the symbol duration allows for a variety of different possible FFT window positions for the evaluation of a symbol. This is indicated in *Fig. 5* for the simple case of synchronization to a single signal. Three possible FFT window positions are indicated as examples. Here, all positions are equivalent with regard to evaluation of the symbol, because all the FFT window positions shown include samples from only one symbol.

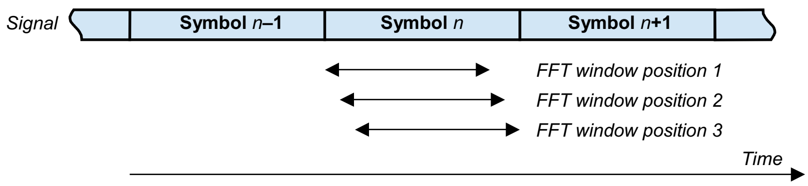

Figure 5
Three possible FFT window positions

The insertion of the guard interval reduces the data capacity because not all of the symbol duration T_s is used for "useful" data.

In a multipath or SFN environment, where many potentially useful signals are available to the receiver, the choice of the FFT window position becomes more complex. A number of different strategies that can be applied are discussed in *Sections 2* and *3*.

All signals with time delays that cannot be absorbed by the guard interval, in the way described above, introduce a degradation of reception, similar to that shown in *Fig. 2*. Any part of a received signal that falls outside the guard interval has an interfering characteristic which is different for T-DAB and DVB-T – due to the different demodulation methods applied (*differential* in the case of T-DAB and *coherent* in DVB-T).

1.2. T-DAB

For T-DAB network planning, the power of all the echoes received within a window of duration Δ (guard interval width) is considered as useful, and contributes positively to the total available signal power. Outside the guard interval, a part of the echo power is associated with the same OFDM symbol as the primary signal, and therefore contributes positively to the total useful signal power.

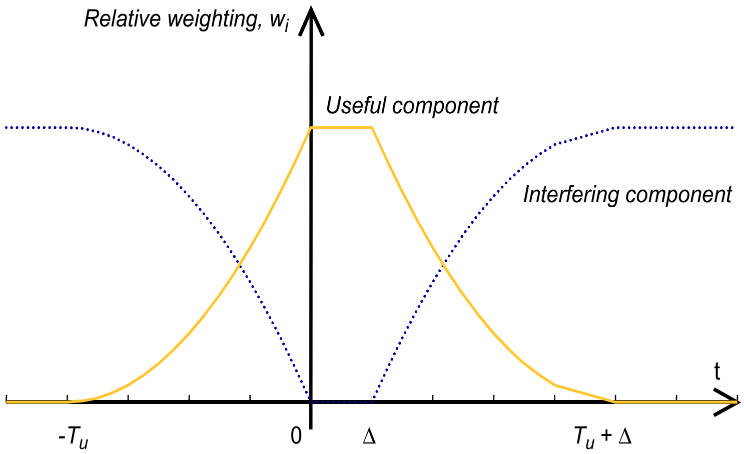

Figure 6
T-DAB model – splitting of the signal power into useful and interfering components

Another part of the echo power is associated with the previous or subsequent OFDM symbol and produces ISI, which has a similar effect to uncorrelated Gaussian noise interference. Therefore, as the echo delay is progressively increased beyond the guard interval, the useful contribution decreases and the ISI increases with a quadratic law. The echo power becomes fully interfering (i.e. it contains no useful power) when the delay is larger than or equal to one OFDM symbol *(see Fig. 6)*.

Mathematically, the rule for splitting the signal power into a useful component and an interfering component is expressed by the *Equation 1*, at the top of the next page.

It must be borne in mind that *I*, the total effective interfering power, is weighted by the established T-DAB-to-T-DAB protection ratio, when being regarded as a source of interference in a coverage calculation.

1.3. DVB-T

In the case of DVB-T, because of the pilot carriers that are needed for coherent demodulation, the total loss of constructive signal components occurs beyond a relative delay of $T_p = T_u/3$. The failure of the equalisation algorithm after T_p (rather than $T_u + Δ$ as in the case of T-DAB) produces a more rapid performance degradation versus the echo delay, than in T-DAB *(see Fig. 7)*. The T-DAB model *(see Fig. 6)* that splits the echo

Equation 1

$$w_i = \begin{cases} 0 & \text{if} \quad t \le -T_u \\ \left(\frac{T_u + t}{T_u}\right)^2 & \text{if} \quad -T_u < t \le 0 \\ 1 & \text{if} \quad 0 < t \le \Delta \\ \left(\frac{(T_u + \Delta) - t}{T_u}\right)^2 & \text{if} \quad \Delta < t \le T_u + \Delta \\ 0 & \text{if} \quad t > T_u + \Delta \end{cases}$$

$$C = \sum_i w_i C_i$$

$$I = \sum_i (1 - w_i) C_i$$

Where:

C_i = the power contribution from the i-th signal at the receiver input
C = the total power of the effective useful signal
I = the total effective interfering power
w_i = the weighting coefficient for the i-th component
T_u = the useful symbol length
Δ = the guard interval length
t = the signal arrival time

power into a useful and an interfering contribution is thus only applicable in the case of DVB-T for echo delays up to T_p, with echoes outside T_p contributing only to the interfering power, independent of their delay.

The corresponding formula is given by *Equation 2*, at the bottom of this page.

Again, as for the T-DAB case, it should be remembered that I, the total effective interfering power, is weighted by the appropriate DVB-T-to-DVB-T protection ratio, when being regarded as a source of interference in a coverage calculation.

A value of $T_u/3$ is regarded as a theoretical limit for T_p and would require an interpolation filter with an infinite number of taps. The formula $T_p = 7T_u/24$ is often quoted and this gives a sensible practical limit in the case of real filter designs. At the present time, many DVB-T receivers do not even reach this performance.

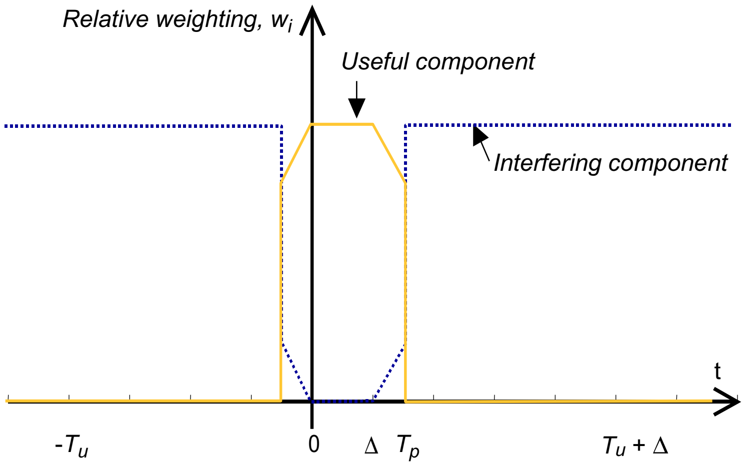

Figure 7
DVB-T model – splitting of the signal power into useful and interfering components.

Equation 2

$$w_i = \begin{cases} 0 & \text{if} \quad t \le \Delta - T_p \\ \left(\frac{T_u + t}{T_u}\right)^2 & \text{if} \quad \Delta - T_p < t \le 0 \\ 1 & \text{if} \quad 0 \le t \le \Delta \\ \left(\frac{(T_u + \Delta) - t}{T_u}\right)^2 & \text{if} \quad \Delta < t \le T_p \\ 0 & \text{if} \quad T_p < t \end{cases}$$

$$C = \sum_i w_i C_i$$

$$I = \sum_i (1 - w_i) C_i$$

Where:

C_i = the power contribution from the i-th signal at the receiver input
C = the total power of the effective useful signal
I = the total effective interfering power
w_i = the weighting coefficient for the i-th component
T_u = the useful symbol length
Δ = the guard interval length
t = the signal arrival time
T_p = the interval during which signals usefully contribute

2. FFT window synchronization

The synchronization of an OFDM receiver is performed in two stages:

- **initial synchronization** in which the receiver is aligned with the symbol rate, and;
- **secondary synchronization** in which the receiver positions the FFT window to demodulate the signal.

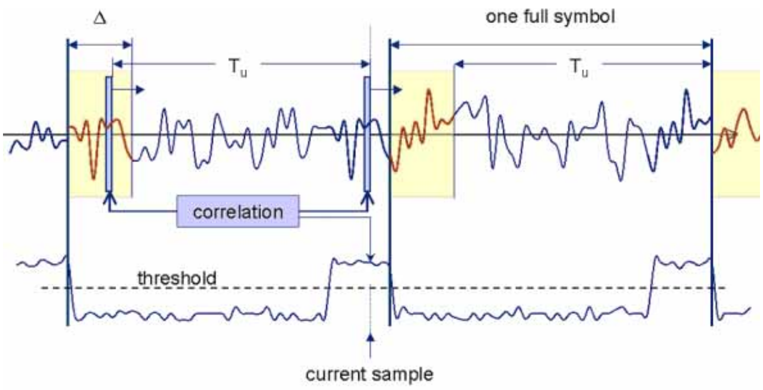

Figure 8
Initial receiver synchronization *(image courtesy of Philips Ltd)*

The initial synchronization is normally done by correlating samples taken T_u apart in time. When the waveform repeats, as shown in *Fig. 8*, the correlator output exceeds a threshold value. From this, the receiver can detect the start of a new symbol period.

In a real multipath environment, the receiver encounters a multitude of echoes which complicates the task of the second-stage synchronization process, i.e. finding the "best" position for the FFT window. As a consequence, various strategies can be applied in order to optimize the receiver performance.

A difference arises from the distinction between direct signals and echoes. In an MFN, where each transmitter acts independently on its own channel/frequency, the receiver may get one direct signal and a number of scattered echoes. The direct signal is not necessarily the strongest signal nor is there necessarily a direct signal at all, particularly in the case of portable or mobile reception. On the other hand, there are also cases where there is only the direct signal present. In an SFN, all transmitters in the network use the same channel/frequency. In this case, the receiver gets a number of direct signals and a number of scattered echoes.

It is necessary to distinguish between different implementations of synchronization strategies in real receivers. These strategies are outlined in *Section 3*.

Most coverage prediction methods use two dimensional prediction models, taking into account only the direct path. Therefore in an MFN, the modelling of the FFT window positioning is simple and unique since there is only one direct path present. In an SFN, receiver synchronization modelling is no longer unique since there are usually several direct-path signals present.

In some three-dimensional prediction models, a multipath propagation environment for each transmitter is considered. Therefore the FFT window positioning for an MFN becomes as complex as that for an SFN.

In planning simulation tools, a natural way to describe the reception situation would be to model real receiver behaviour. Unfortunately, the receiver FFT window positioning is not prescribed in detail in either the T-DAB or DVB-T system specifications. This means that all manufacturers have their own solutions and, moreover, they regard these various solutions as confidential – making a single description of receiver FFT window positioning difficult.

A further difference arises from the fact that real receivers have to account for the time variation of the transmission channel, whereas software modelling of the receiver FFT window positioning usually assumes a static reception situation. (This, to some extent, is justified by the different time scales of successive synchronization instants and the time variation of shadow fading in a transmission channel.) This means that a real receiver will not show exactly the same synchronization behaviour as that described in the simple model cases below. Real receiver synchronization has to deal with all existing echoes in the multipath environment, while receiver modelling usually does not.

3. Synchronization strategies

3.1. General

This section describes five different strategies for second-stage synchronization (i.e. positioning of the FFT window) that are commonly used in receiver modelling. Four of them are relatively simple and straightforward strategies, while the fifth is an idealised optimal strategy.

The strategy employed by a receiver determines:

- which peak, in the time-domain impulse response of the received signal, the receiver uses for synchronization;
- where the receiver sets the FFT window relative to this peak.

In a single-signal environment, the synchronization configuration is simple and clear. The principle was already explained in *Section 1* and can be seen from, for example, *Fig. 4*.

FFT window synchronization is of particular importance for mobile and portable reception, when the receiver will need to be able to synchronize in a rapidly changing environment and in the presence of pre- and post-echoes.

3.2. Strongest signal

A natural approach for the FFT window positioning is to synchronize to the strongest signal, in a similar way to that shown in *Fig. 4* for a single signal. In order to demonstrate the principle, a configuration with four signals is chosen as an example. *Fig. 9* shows the channel response function for the configuration, where the peaks represent a characteristic time instant of the signals, such as the start of symbol *n*.

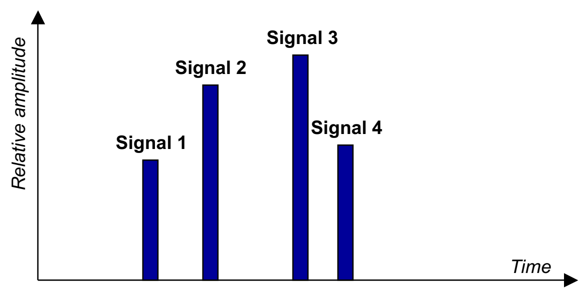

Figure 9
Synchronization to the strongest signal (signal 3); impulse response

Signal 3 is the strongest signal. Accordingly, the FFT window is synchronized to signal 3. Since relevant contributions of further signals may be found preceding signal 3 or following signal 3, it seems reasonable to locate the centre of the FFT window at the centre of symbol *n* of signal 3. This is depicted in *Fig. 10*. In the example, signals 3 and 4 contribute fully to the evaluation of symbol *n*, whereas the FFT window exhibits an overlap with symbol *n*+1 of the signals 1 and 2, which results in a certain amount of ISI.

A more sophisticated synchronization strategy, based on the strongest signal approach, would not be fixed to the centre of the symbol duration but would check for better positions within the symbol duration of the strongest signal. In the chosen example, it would be advantageous to move the FFT window a tiny bit backwards in time to avoid the small amount of ISI arising from the overlap with symbol *n*+1 of signal 2. Also, the inter-symbol interference from signal 1 would be reduced.

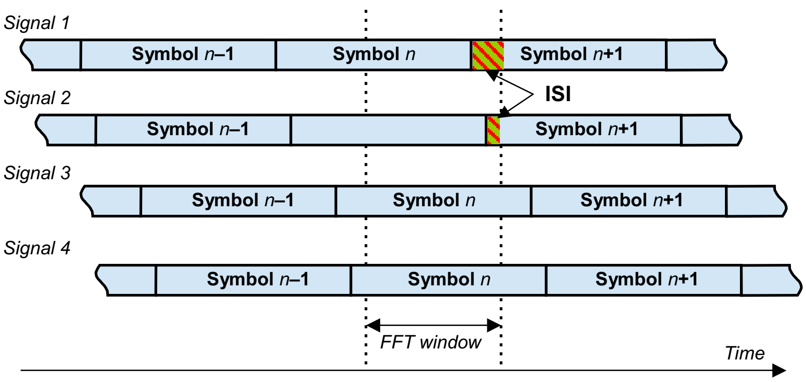

Figure 10
Synchronization to the strongest signal (signal 3); FFT window position

3.3. First signal above a threshold level

This strategy takes the first signal of the time impulse response as a reference for the FFT window. Normally, a minimum threshold level is necessary for a signal in order to be accepted as a trigger. Again the 4-signal configuration of the previous section is taken as an example. The impulse response is given in *Fig. 11* with the threshold value indicated by a horizontal dashed line.

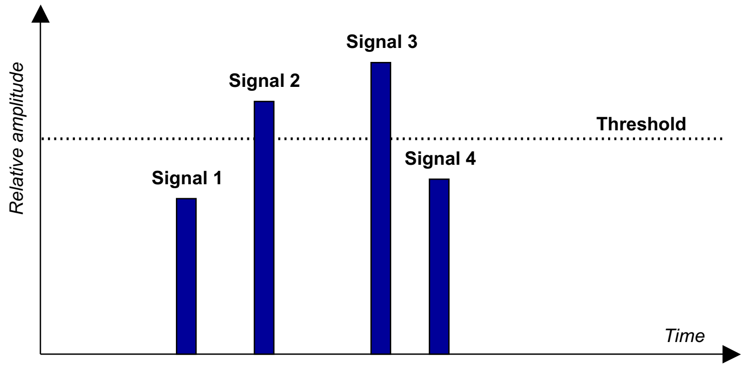

Figure 11
Synchronization to the first signal above a threshold (signal 2); impulse response

The first signal above the threshold is signal 2. It serves here as the trigger for the FFT window. If the threshold is chosen reasonably, it can be expected that there is no significant signal preceding signal 2 and, therefore, it is logical to align the end of the FFT window with the end of the symbol n of signal 2. This is indicated in *Fig. 12*.

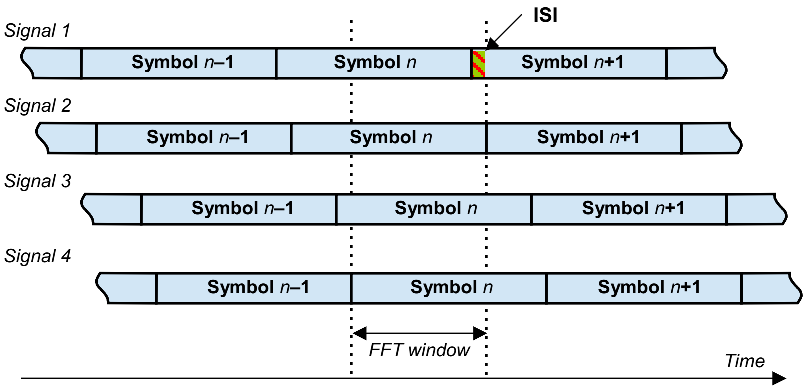

Figure 12
Synchronization to the first signal above a threshold (signal 2); FFT window position

With this synchronization strategy, signals 2, 3 and 4 contribute fully constructively, whereas signal 1 adds a certain amount of ISI.

The choice of the threshold value is a specific issue of this synchronization strategy. It may be taken as the power corresponding to the minimum field strength or, more pragmatically, as a value, say 6 to 10 dB, below the strongest signal.

In a recent workshop, EICTA has indicated that various manufacturers apply the "first signal above a threshold level", or a strategy similar to it.

3.4. Centre of gravity

In this case the receiver looks at the impulse response, calculates the "centre of gravity" of the impulse response spectrum and centres the FFT window on that point in time:

$$t_c = \frac{\sum_i p_i t_i}{\sum_i p_i}$$

where

t_c = centre of gravity

p_i = power of the i-th signal of the impulse response

t_i = time of the i-th signal of the impulse response

The impulse response of the chosen example, with the corresponding centre of gravity indicated by a dashed line, is given in *Fig. 13*.

In this example, signals 2 and 3 fully contribute constructively. Signals 1 and 4 show a small amount of inter-symbol interference arising from an overlap of the FFT window with symbol n+1 of signal 1 and with symbol n–1 of signal 4. This is depicted in *Fig. 14*.

The centre of gravity approach responds well to pre-echoes and delayed signals of similar amplitude, since it does not fix the FFT window to a particular signal but takes into account the average behaviour of the impulse response of the transmission channel. On the other hand, it can lead to ISI in cases where other strategies may not lead to ISI: for example, most two-echo cases, separated by virtually the whole guard interval, would cause difficulties for this strategy unless the two echoes were of equal power.

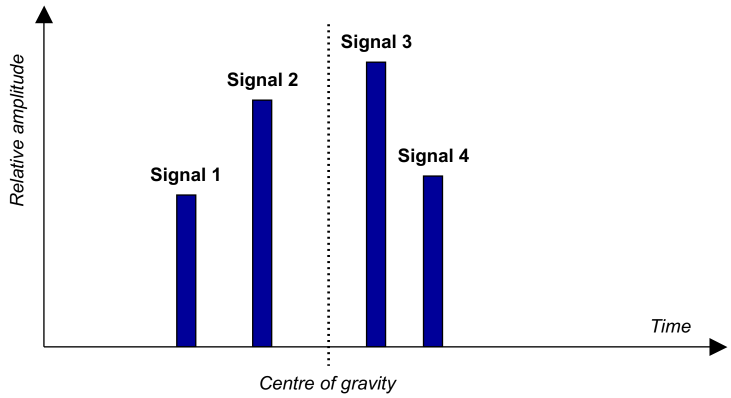

Figure13
Synchronization to the centre of gravity (between signal 2 and 3); impulse response

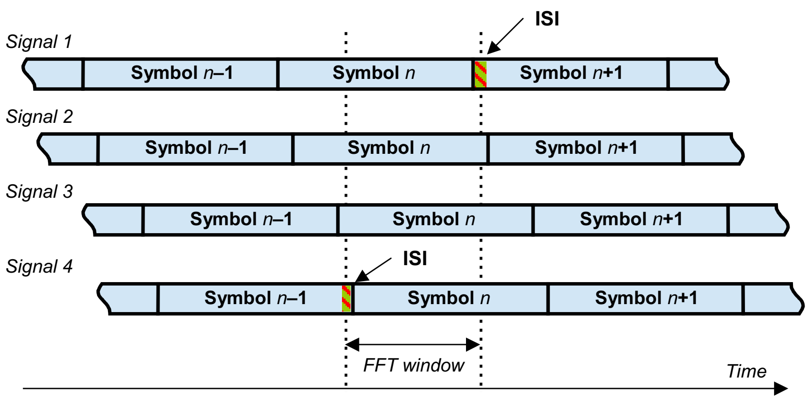

Figure 14
Synchronization to the centre of gravity (between signal 2 and 3); FFT window position

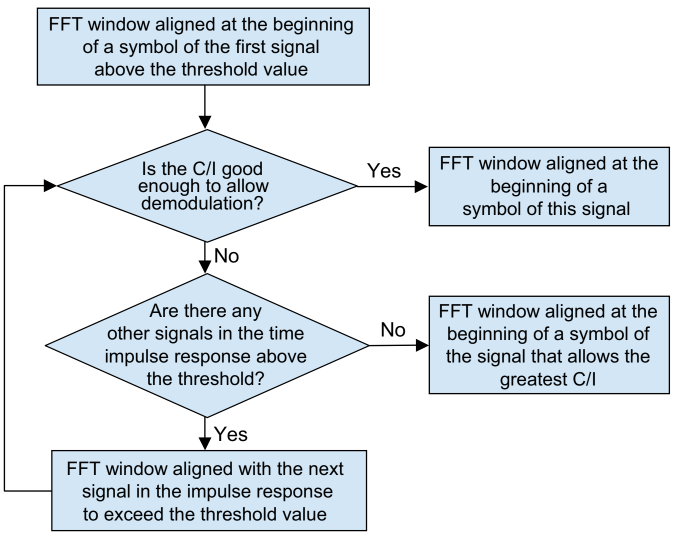

Figure 15
Flowchart describing the quasi-optimal strategy

3.5. Quasi-optimal

This strategy builds on that described in *Section 3.3.*, in an attempt to approach the " Maximum C/I " described in the next section.

The first signal of the impulse response above a minimum threshold level is taken as a reference for the FFT window. The process is described in the flowchart shown in *Fig. 15*.

3.6. Maximum C/I

Whereas the previously-discussed strategies all offer means of quickly finding a good FFT window position, an optimal choice would be a position where the effective C/I is maximized. This position, however, is not easily found and would in general take too much time to be calculated. Therefore, normally one of the above simpler strategies, or a combination of them, is applied.

Such simpler approaches can be justified by the fact that the optimum C/I will often show a relatively flat maximum, i.e. errors introduced by sub-optimal synchronization are small. But there are also difficult configurations possible, e.g. in a two-echo case, if the difference in delay is close to the guard interval, there is only one position that will result in no ISI, so the optimum here would be very sharp.

Note that the method described in *Section 3.5.* does not attempt to find a position for the FFT window that gives the best C/I. It merely seeks to find a position for the window at which the C/I is good enough to allow demodulation and decoding with an acceptable error rate.

Receiver manufacturers indicate that the evaluation of C/I is by no means trivial for a DVB-T receiver and, for a DVB-T mode with a large guard interval of $T_u/4$, there seem to exist theoretical limits for the evaluation of C/I which would prevent the application of a " maximum C/I " synchronization strategy in this case.

With regard to receiver modelling in computer simulations (e.g. for coverage calculations), the detection of the maximum C/I position of the FFT window is not a major problem. A simple but time-consuming approach would be to scan the time period of interest with an appropriate step size, calculate the C/I for each sampling point and use the time position with the maximum C/I as the reference.

A more sophisticated strategy to find the maximum C/I position is based on the observation that the maximum C/I is always found at a position where the FFT window is aligned with the start, or the end, of one of the incoming signals for the symbol under consideration. A check of all these possible positions, which amounts to 2N evaluations of C/I for N signals, then gives the maximum C/I position. Experience with practical software implementations shows that the computational effort is about twice that of the basic strategies described in *Sections 3.2.* to *3.4.*.

4. Studies

4.1. Theoretical regular T-DAB network

As described above, the FFT window synchronization strategy has an effect on the receiver performance in the presence of more than one signal (coming from either one transmitter or from transmitters in an SFN). In order to demonstrate how FFT window synchronization influences the receiver performance, a regular 7-transmitter hexagonal network, having characteristics similar to that of the Wiesbaden 95 Band III T-DAB reference network, has been used. The network is open, and the central transmitter can be given a time offset. T-DAB Mode I ($\Delta = 246$ µs) is used in all the simulations shown in this example. With no time offset on the central transmitter, the area inside the hexagon is virtually free from inter-symbol interference. By applying a time offset of 300 µs to the central transmitter, a situation where ISI occurs has been deliberately created.

It must be emphasized that this theoretical network is used solely for the comparison of receiver synchronization methods and is not intended to represent a network that would actually be constructed. However, in real networks, time delays do occur.

A study carried out in Germany, to investigate the above synchronization modes in this theoretical network, indicates that the synchronization strategy is a crucial feature for the performance of a receiver.

The results for four synchronization strategies are given in *Figs 16 - 19*.

The results show that, for very critical reception situations, the difference between the predicted coverage probabilities for the various synchronization strategies may amount to as much as 20%. For the high probability range, the differences are found to be between 5 and 10%.

Moreover, the investigation showed that knowledge about the synchronization behaviour of the receivers is important for reliable T-DAB network planning.

4.2. Real network simulation

In a case study in Switzerland, a realistic DVB-T SFN simulation in the Zürich region was performed. The network consisted of three transmitters, the characteristics of which were optimized in terms of ERP and time

Abbreviations

C/I	Carrier-to-Interference ratio
CEPT	European Conference of Postal and Telecommunications Administrations
DAB	Digital Audio Broadcasting (Eureka-147)
DVB	Digital Video Broadcasting
DVB-T	DVB - Terrestrial
EICTA	European Information, Communications and Consumer Electronics Technology Industry Associations
ERP	Effective Radiated Power
FFT	Fast Fourier Transform
IFFT	Inverse Fast Fourier Transform
ISI	Inter-Symbol Interference
MFN	Multi-Frequency Network
OFDM	Orthogonal Frequency Division Multiplex
SFN	Single-Frequency Network
T-DAB	Terrestrial - Digital Audio Broadcasting

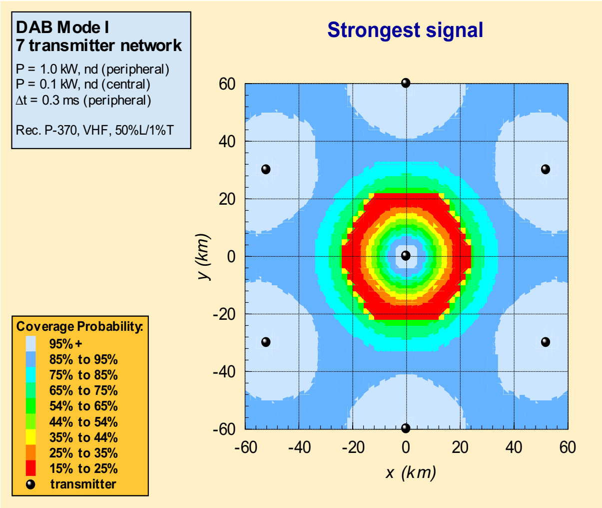

Figure 16
Synchronization – " Strongest signal "

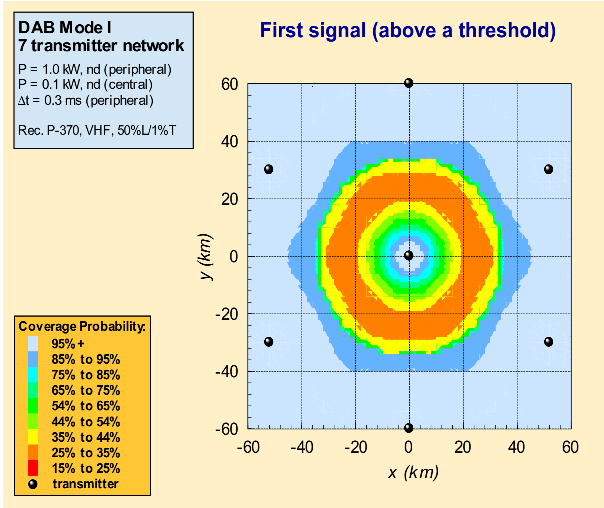

Figure 17
Synchronization – " First Signal above a threshold "

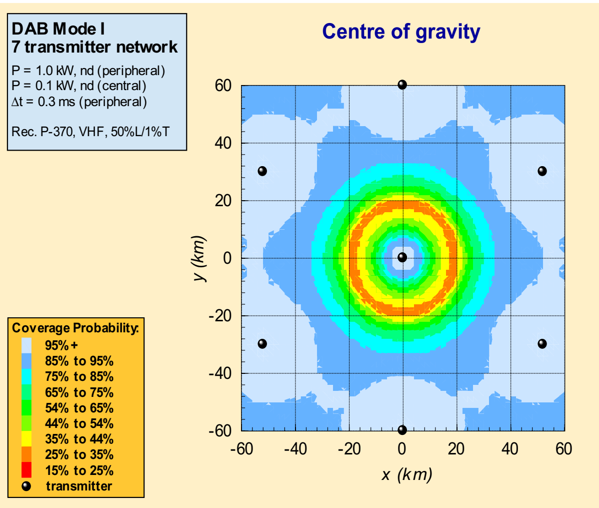

Figure 18
Synchronization – " Centre of gravity "

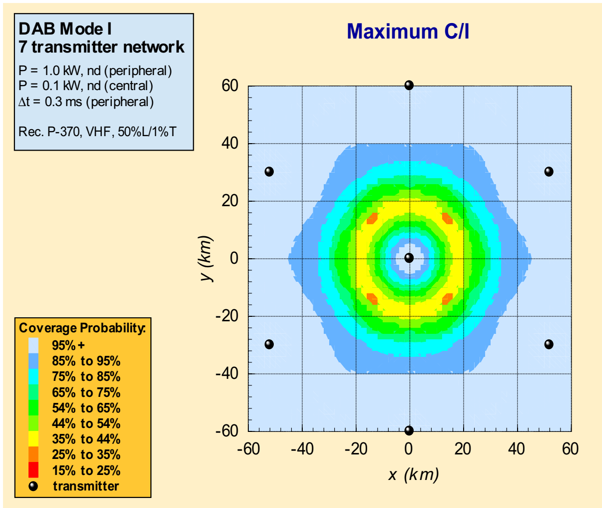

Figure 19
Synchronization – " Maximum C/I "

delays in order to maximize coverage. Different receiver synchronization algorithms were compared and the resulting population coverage was calculated. *Table 1* gives the results. The trigger threshold used in simulations for the first signal strategy was 20 dB above the noise level. However, it should be mentioned that only signals inside the guard interval were regarded as wanted, whereas signals outside the guard interval were treated as interference (cliff-edge transition), which is an approximation to the response in *Fig. 7*.

Table 1
Percentage of the population covered by a digital SFN (16-QAM $^3/_4$, $\Delta = {}^1/_8$) in the Zürich region for different synchronization algorithms at the receiver. ERP of all transmitters is 40 dBW.

FFT Window	Fixed outdoor reception		Portable indoor reception	
	Band III	Band IV	Band III	Band IV
Start = the strongest signal	85.8 %	86.1 %	78.0 %	70.7 %
Start = the first signal over threshold	97.1 %	95.3 %	83.8 %	74.2 %
Centre = the centre of gravity of all signals	77.8 %	77.7 %	63.7 %	59.9 %

The results show that the choice of the synchronization strategy is crucial for the coverage of the network. The study concluded that a strategy that aligns the FFT window with the start of the first signal above a certain threshold is the best choice of the three considered.

4.3. Measurements on DVB-T receivers

Since manufacturers' information is not available, laboratory measurements in Italy have been performed in order to investigate the synchronization strategies that are employed by nine receivers.

The study came to the conclusion that the majority of the available DVB-T receivers uses the "threshold" algorithm (but with different threshold levels). Only one receiver (based on the 2k mode and therefore not directly suitable for wide SFNs) adopts the "centre of gravity" method. No receiver was identified using the " maximum C/I " algorithm.

5. Conclusions

As the T-DAB and DVB-T systems differ significantly in their specifications, T-DAB receivers work differently from DVB-T receivers in some crucial aspects. In particular, T-DAB receivers do not exhibit the time cut-off that is a feature of all DVB-T receivers, so longer echoes can be used constructively in T-DAB systems.

Different possible strategies of FFT window synchronization in OFDM receivers have been reviewed. Since the FFT window-positioning strategy may significantly affect the coverage probability in a network:

1) it is important for network planners to know about the synchronization strategies adopted by current and future receivers, and;
2) it should be of interest to the receiver manufacturers to know which strategy has been assumed for the planning of OFDM networks.

These considerations should seriously be taken into account by the manufacturers (EICTA) and by the international and national bodies responsible for frequency planning (CEPT and member administrations).

In a recent meeting between the EBU and EICTA, it was indicated that the "first signal above a threshold" strategy is most commonly used in receivers. This agrees with the findings of the study mentioned in *Section*

Dr Roland Brugger graduated from the Ludwig-Maximilians University, Munich, in 1981 and obtained a PhD in physics in 1987. In the same year, he joined the Frequency Management Section of the Institut für Rundfunktechnik (IRT) in Munich, where he worked on statistical aspects and on combinatorial optimisation problems in the frequency planning field. Since 1992, he has been involved in frequency and network planning for digital terrestrial broadcasting services.

Roland Brugger participated in the T-DAB and DVB-T Planning Meetings in Wiesbaden (1995), Chester (1997) and Maastricht (2002), and is presently engaged in the preparation of the RRC04/06. He is a member of ITU-R TG6/8 and of the EBU project groups B/CP-R and B/CP-T.

David Hemingway joined the BBC's Transmission department in 1988 and worked on HF, VHF and UHF transmitting stations before becoming responsible for microwave link projects in the mid-1990s. After the privatisation of BBC Transmission in 1997, he helped to establish the Spectrum Planning group within Crown Castle UK Ltd.

In 2000, Mr Hemingway moved to work for Nokia, where he managed a European group of 3G radio planners and worked in Sweden and Switzerland as well as the UK in that capacity. In 2001, he rejoined the BBC and now works as a Senior R&D Engineer in the Spectrum Planning Group at Kingswood Warren. He is currently one of the leaders of the BBC's preparatory work for the RRC in 2004/06.

4.3., and corresponds to the technique assumed by most network planners when designing their networks. Note however that the value of the threshold used is not specified.

Acknowledgement

This article results from work done in EBU Project Group B/CP-T (Conference Preparations – Television). The authors would like to thank all the colleagues who participated in the preparation of the article.

The article is dedicated to Jørn Andersen.

25 July 2003