

Performance Analysis and Receiver Architectures of DCF77 Radio-Controlled Clocks

Daniel Engeler

Abstract—DCF77 is a longwave radio transmitter located in Germany. Atomic clocks generate a 77.5 kHz carrier which is amplitude- and phase-modulated to broadcast the official time. The signal is used by industrial and consumer radio-controlled clocks.

DCF77 faces competition from the Global Positioning System (GPS) which provides higher accuracy time. Still, DCF77 and other longwave time services worldwide remain popular because they allow indoor reception at lower cost, lower power, and sufficient accuracy.

Indoor longwave reception is challenged by signal attenuation and electromagnetic interference from an increasing number of devices, particularly switched-mode power supplies.

This paper introduces new receiver architectures and compares them with existing detectors and time decoders. Simulations and analytical calculations characterize the performance in terms of bit error rate and decoding probability, depending on input noise and narrowband interference.

The most promising detector with maximum-likelihood time decoder displays the time in less than 60 s after powerup and at a noise level of $E_b/N_0 = 2.7$ dB, an improvement of 20 dB over previous receivers.

An FPGA-based demonstration receiver built for the purposes of this paper confirms the capabilities of these new algorithms. The findings of this paper enable future high-performance DCF77 receivers and further study of indoor longwave reception.

I. INTRODUCTION

Most clocks and watches require periodic adjustment, either to correct movement imprecision, to switch between summer and winter time, or after a battery replacement. Radio-controlled clocks usually run on a quartz oscillator and synchronise periodically to a radio time signal. This automatic adjustment coined the marketing name of “atomic watch” and led to the commercial success of such wristwatches, alarm clocks, but also industrial receivers which require reliable and accurate time.

Time distribution through radio comes in several forms and frequencies: Spoken time announcements on news channels, FM RDS (radio data system), GPS (global positioning system), and transmitters on longwave and shortwave. Each system has its advantages: GPS provides worldwide time with nanosecond accuracy, but only with a relatively complex receiver and a line-of-sight to a satellite. Longwave receivers on the other

hand are simpler and work indoors. Due to their fundamentally different architectures, GPS and longwave are ideal backup time systems for each other.

The longwave transmitter DCF77 lowers its carrier amplitude once per second for either 100 ms or 200 ms, which the receiver detects as a bit 0 or 1. The resulting 60 bits per minute contain a minute start signal, the current minute, hour, and date [1], as shown in Table II on the last page. The current second is derived by counting seconds after the minute start.

This amplitude modulation (AM) enables millisecond accuracy [2]. DCF77 also uses phase modulation (PM) [3], [4] which enables accuracies of a few microseconds. Fig. 1 shows the modulation symbols. The DCF77 AM and PM are orthogonal, such that a receiver can detect either or both modulations. Other popular longwave time transmitters such as WWVB (U.S.), MSF (UK) and JJY (Japan) use mostly AM with different encodings.

Instruction manuals of radio-controlled clocks recommend to place the device near a window and away from electrical devices, both in attempt to improve the signal-to-noise ratio. Experience shows that switched-mode power supplies are a major threat to indoor longwave reception. This paper analyses the performance of DCF77 receivers in the presence of such noise and interference and proposes new algorithms with higher performance than what is currently available.

II. DETECTOR TYPES

A DCF77 receiver consists of a detector and a time decoder (Fig. 2). The detector receives the radio signal, demodulates and synchronises to it, then converts it into a stream of 1 bit/s. The following sections introduce several detector types.

A. Diode detector (existing)

The diode detector (Fig. 3) is the simplest and probably the most frequently used amplitude detector. It can be built with discrete components which include a diode for rectification, hence its name. It uses the following structure:

- 1) After the antenna amplifier, bandpass filter H_1 removes unwanted signals outside of the DCF77 frequency band.
- 2) The signal is rectified by computing either $|x|$ or x^2 (the latter is preferred for fewer harmonics).
- 3) A lowpass filter H_2 removes the high frequency components created by the rectification and leaves only the signal envelope. The frequency separation is twice the

D. Engeler (e-mail: engeler@alumni.ethz.ch) is with Zühlke Engineering AG, 8952 Schlieren, Switzerland, which funded the production of the Demonstration Receiver.

This is a draft version of the paper which was published in the IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control, May 2012, vol. 59, no. 5, pp. 869–884, DOI <http://dx.doi.org/10.1109/TUFFC.2012.2272>




Fig. 1. DCF77 symbols: One second of the DCF77 signal contains an AM bit (AM0 or AM1, shown here after demodulation), and a PM bit (PM0 or PM1). Also shown are the derived symbols $AM^{1/2} = (AM0 + AM1)/2$, and $AM\Delta = AM1 - AM0$. When correlating with these symbols, a receiver must adapt the symbol shape to the frequency response of its input stage.




Fig. 2. General DCF77 receiver structure.




Fig. 3. Diode detector with optional H_3 for averaging.

carrier bandwidth, therefore a simple first-order filter is sufficient.

- 4) The AM envelope falling edge (which indicates the start of a second) is found with a threshold relative to the high level.
- 5) The AM bit is sampled at 0.15s after the start of a second.

Choosing the bandwidth B_1 of H_1 requires a trade-off between immunity (against noise and interference) and precision. For the simulations described in section IV, $B_1 = 40$ Hz (immune, but not precise) and $B_1 = 2583$ Hz (precise, but not immune) are used. The latter contains the main lobe and the first side lobes of the DCF77 signal and is the maximum useful bandwidth for DCF77 detection.

The diode detector can be interpreted as a special case of a quadrature detector with a mixing frequency of 0 Hz [5]. Simulations show that quadrature demodulation without carrier synchronisation offers no advantage over the diode detector.

If B_1 is chosen rather large, either for improved precision or reduced filter complexity, an additional averaging filter H_3 can be used to improve immunity. The averaging window (Fig. 4) starts when the rising edge of AM bit 0 has settled at 0.1 s + τ until 0.2 s, where $\tau \propto 1/B_1$. With $N = \lfloor (0.1 \text{ s} - \tau) f_s \rfloor$ samples being averaged, the filter response is $H_3(f) = \frac{\sin \pi f N}{N \sin \pi f}$.

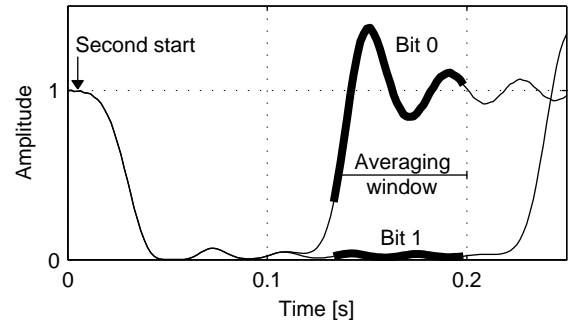


Fig. 4. Example AM envelopes (after lowpass filter H_2) of diode detector with averaging, using a sharp bandpass filter H_1 with $B_1 = 60$ Hz.




Fig. 5. IIR (infinite impulse response) filter implementation of the Goertzel algorithm, with infinite window length N [6]. Only delays, additions, and constant multiplications are used, enabling efficient hardware implementations.

B. Goertzel detector (new)

The Goertzel detector is a new DCF77 detector introduced by this paper. It is named after the well-known Goertzel algorithm which computes the k -th value of an N -point DFT (discrete Fourier transform):

$$X(k) = \sum_{n=0}^{N-1} x(n)e^{-j2\pi nk/N}. \quad (1)$$

If this is applied at the DCF77 carrier frequency, the complex value X contains the amplitude and phase of the carrier, from which the AM and PM can be derived.

The window length N of a DFT determines its frequency resolution. This can be used to adjust a detector's selectivity to the desired signal. For infinite N (wideband), the algorithm can be implemented efficiently as shown in Fig. 5. For finite N , [6] describes a sliding DFT using N memory cells. In this case, the selectivity is limited by receiver memory.

The Goertzel detector introduced by this paper combines the advantages of these two algorithms: Variable selectivity with a low (and fixed) amount of memory. Starting with the infinite- N implementation of Fig. 5, the two memory cells are scaled periodically. This scaling results in an exponentially weighted moving average which has the effect of a time-domain window and thus defines the frequency selectivity.

If the scaling is performed every carrier cycle, the scaling constant k corresponds to a 3 dB bandwidth of approx. $0.32 \times (1 - k)f_c$. A time domain interpretation of $k = \alpha^{1/n}$ is that a

carrier amplitude step reaches $1 - \alpha$ of the final value within n cycles.

The DCF77 modulations AM and PM are orthogonal in two aspects:

- 1) AM uses only the amplitude, PM uses only the phase.
- 2) AM is used only between 0 s and 0.2 s after the start of a second, whereas PM is used between 0.2 s and 1 s.

It is therefore obvious that the AM and PM components should be detected separately, each with the best possible frequency response. A similar approach is described in [7]. The optimum frequency response would be a matched filter which is difficult to realise. For the Goertzel algorithm used here, the scaling constants are adjusted such as to fit the spectrum of the AM or PM component as good as possible, resulting in 3 dB bandwidths of 15 Hz (AM) and 930 Hz (PM).

Fig. 6 shows the block diagram of the Goertzel detector. Three instances of the Goertzel algorithm detect the carrier, AM, and PM components, resulting in the complex values X_{Carrier} , X_{AM} , and X_{PM} . Fig. 7 shows the resulting frequency responses.

The carrier phase relative to the receiver is $\arg(X_{\text{Carrier}})$, which can be calculated with CORDIC or a similar algorithm. The Goertzel scaling constant for the carrier component depends on the maximum clock frequency error of the receiver. The smaller this error, the better the selectivity.

The AM envelope can be calculated as $|X_{\text{AM}}|$. When using CORDIC, vector rotation with the carrier phase simplifies this to $\text{Re}(X_{\text{AM}})$.

Similarly, the PM phase can be calculated as $\arg(X_{\text{PM}} - X_{\text{Carrier}})$. Since the phase deviation from the carrier is only $\pm 13^\circ$ [1], after vector rotation this can be approximated by the imaginary part.

Detecting AM and PM separately results in 1 bit/s for AM and 1 bit/s for PM. These bits are equal only in seconds 15–58 [1], therefore they can be combined by the time decoder once the minute start is known.

The Goertzel detector is suited for efficient fixed-point implementation on microcontrollers and FPGAs. For the purposes of this paper, it was implemented on the FPGA-based Demonstration Receiver described in section VI.

C. CIC detector (new)

The CIC detector is a new DCF77 detector introduced by this paper, together with the Goertzel detector described above. The CIC detector is named after the cascaded integrator-comb filter, based on the following motivation:

The DCF77 carrier frequency of 77.5 kHz is 30 times as large as its bandwidth (2583 Hz for the main and first side lobes), therefore mixing to a lower frequency and down-sampling simplifies processing.

Fig. 8 shows the quadrature zero-IF (intermediate frequency) CIC detector. The antenna signal is first quadrature-mixed to DC. The I and Q components are then lowpass-filtered and decimated by several CIC stages.

Similar to the Goertzel detector, the signal components PM, AM, and carrier are filtered individually to each component's optimum frequency response. Where the Goertzel detector

applied 3 different instances of the Goertzel algorithm in parallel, the CIC detector stages operate sequentially. Each CIC stage therefore benefits from the previous stage, which reduces the computational workload.

The CIC filter was chosen because it is a hardware-efficient decimation filter, requiring only delays and additions (no multiplications). CIC filters are therefore well-suited for implementation on FPGAs.

CIC filters have linear phase, which is a useful property for a time receiver. The frequency-independent delay results in an undistorted signal whose constant delay can be easily compensated. This is important mainly for wideband PM reception, where a non-linear phase filter would degrade the accuracy of the received time.

The CIC filter design must trade off aliasing, precision, stopband attenuation and passband flatness. Fig. 7 shows the filter responses which are useful for DCF77 reception, and which were used in the simulation model for this paper.

D. Matched filter detector (benchmark)

In the matched filter detector, a matched filter correlates the antenna signal with a bank of known reference signals (Fig. 9). The bank with the highest correlation is selected as the received signal.

This detector is optimal in the presence of AWGN (additive white Gaussian noise) and therefore serves as a benchmark. It is computationally intensive and here only used for simulation.

The DCF77 AM and PM are orthogonal (as described above) and are therefore detected separately. The AM filter uses only the part of the signal between 0 s and 0.2 s, the PM filter between 0.2 s and 1 s.

To be optimal, a matched filter requires knowledge about the transmitter, propagation, and receiver. This is possible for DCF77 since all components are well-defined.

III. DETECTOR SYNCHRONISATION

A DCF77 receiver must synchronise to various parts of the DCF77 signal, in the following order:

- 1) Carrier frequency: This is necessary for demodulation. Usually a receiver has a quartz-based reference frequency from which it approximates the 77.5 kHz carrier frequency.
- 2) Start of second: One second of the DCF77 signal contains an AM bit and a PM bit. These bits can be detected only if the receiver knows the start of a second.
- 3) Start of minute: One minute of the DCF77 signal contains the full time and date. To decode this, the receiver must know in which second the minute starts. This is handled by the time decoder described in section V.

This section describes step 2 in more detail.

The simplest synchronisation method uses the AM falling edge: Averaging the AM envelope over several seconds results in the envelope high level. A threshold relative to the high level detects the falling edge. Using this method, precisions of up to 50 μs have been achieved, but only with a large input bandwidth which makes the receiver susceptible to noise [1], [4].

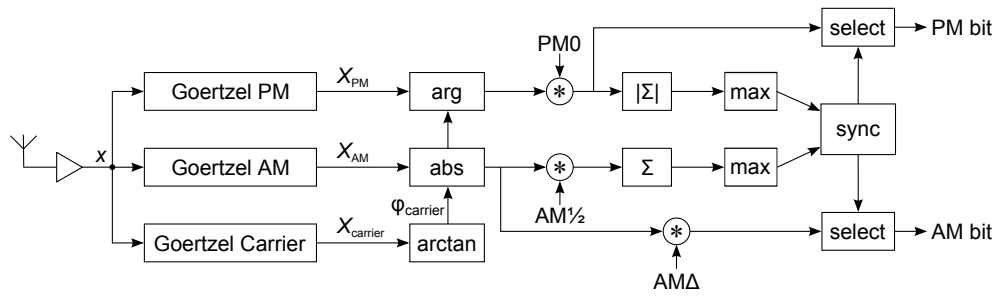


Fig. 6. Goertzel detector with single-second synchronisation. $AM^{1/2}$ is the envelope average of AM bits 0 and 1 (Fig. 1). $AM\Delta$ is the difference envelope. PM0 is the 512-bit pseudo-random phase pattern.

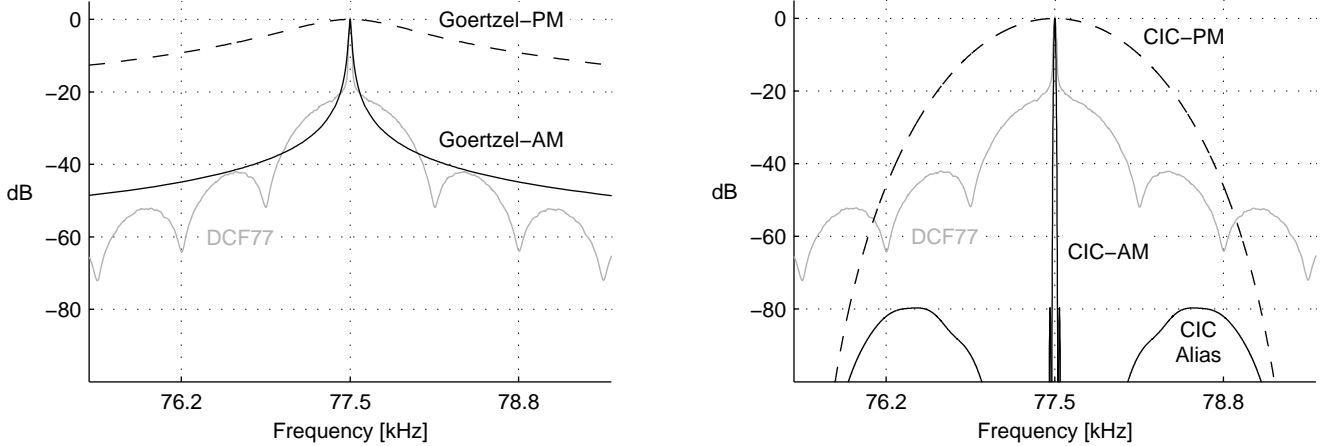


Fig. 7. Frequency responses of the Goertzel and CIC detectors, showing the different selectivities for optimum AM and PM detection. The DCF77 spectrum is shown in light grey, with clearly visible side lobes. Outside of the DCF77 spectrum (main and first side lobes), the CIC detector has much higher suppression than the Goertzel detector.

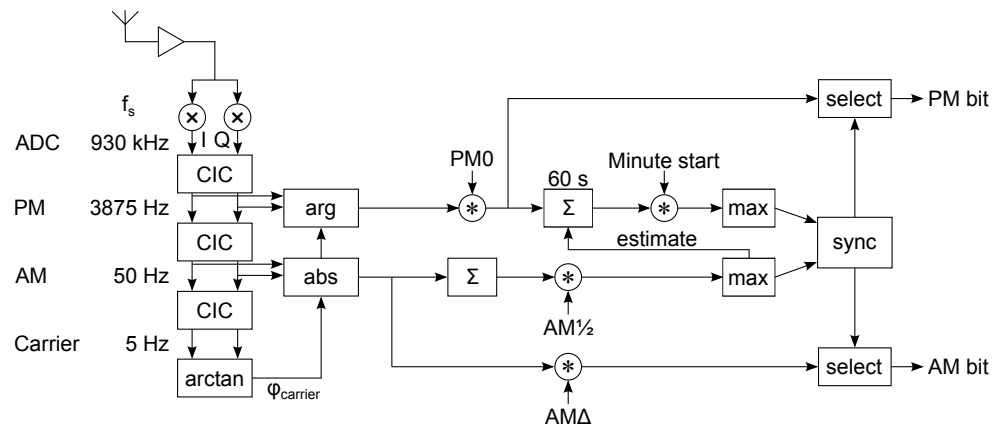


Fig. 8. CIC detector with full-minute synchronisation. This is the highest-performing DCF77 detector introduced by this paper. The sampling frequencies f_s shown were used for the simulations.

Another AM synchronisation method (as used for the simulations in this paper) correlates the received AM envelope with $AM^{1/2}$, the average envelope of AM bits 0 and 1 (Fig. 1). This results in a triangular correlation peak which is approx. 400 ms wide. For all except the last second of a minute, the correlation accumulates linearly which averages noise. In principle, any amount of noise can be rejected by waiting long enough.

For PM synchronisation, the demodulated phase difference to the carrier is correlated with PM0, the 512-bit pseudo-

random code. This enables precisions of up to $13 \mu s$, approx. one carrier cycle [1].

The PM correlation peak is positive for PM0 (a second with PM bit 0) and negative with equal amplitude for PM1, therefore simple addition over several seconds would cancel the correlation peak. This can be solved using two methods:

- 1) The absolute value of the correlation is accumulated (also described by [7]). However, this prevents noise averaging which limits performance.




Fig. 9. Matched filter detector. This detector type is optimal in the presence of AWGN and therefore serves as a simulation benchmark.

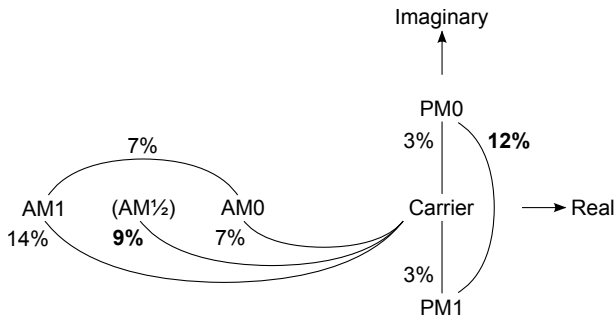


Fig. 10. Quasi-constellation diagram of DCF77 modulation showing energy of signal difference relative to the carrier. The energy difference from the carrier is biggest for AM (9%), which provides the best second synchronisation. The energy difference between bit 0 and bit 1 is biggest for PM (12%), which provides the best second detection.

- 2) The correlation is performed over a full minute using the minute start pattern in PM bits 0 to 14 [1]. This results in not only the second synchronisation, but also the minute synchronisation. Simulations show that approx. 5 dB more noise can be handled than with the previous method for equal reception duration.

Synchronisation performance depends on the symbol constellation. The DCF77 symbols AM0, AM1, PM0, and PM1 cannot be simplified to a single point in the complex plane, but Fig. 10 attempts to visualise the constellation nonetheless.

AM and PM synchronisation result in identical second start positions (though with different precisions). Fig. 10 and simulations show that for AWGN, AM synchronisation is more robust than PM synchronisation. Therefore AM synchronisation can be used to estimate the initial position of the PM correlation, which reduces memory requirements. It is expected that approx. 8000 bytes of memory are sufficient for full-minute AM + PM synchronisation with a resolution of 3875 Hz. Such a combined algorithm enables high-immunity low-cost receivers with reasonable accuracy.

IV. DETECTOR PERFORMANCE

The performance of a DCF77 receiver can be quantified in various ways, such as:

- Maximum input noise
- Minimum input field strength

- Accuracy
- Directional selectivity
- Probability of incorrect decoding
- Time until first decoding
- Power consumption

A receiver consists of a detector and a decoder which operate sequentially, therefore their performance can be analysed independently. After having introduced several detector types, this section looks at detector performance, quantified as output BER (bit error rate) depending on input noise and interference. The BER is based on 1 bit/s, assuming identical AM and PM bits. For zero noise, any detector should have BER = 0. With increasing noise, the BER also increases, finally reaching BER = 0.5 where decoding is not possible anymore.

An accurate analytical representation of the DCF77 signal, noise, interference, and detector is difficult, therefore simulations are used instead. For each detector described in the previous sections, a simulation model was written.

A. Additive white Gaussian noise

To investigate the impact of AWGN on detector performance, the following simulations were run:

- 1) One second of the DCF77 signal with a random bit 0 or 1 is generated.
- 2) AWGN with a specified power density is added. See Fig. 11 for example noise levels in the time- and frequency domain.
- 3) This signal is fed to a detector model which calculates the bit.
- 4) The resulting bit is compared with the original bit. If the bits differ, a bit error has occurred.

For each detector type and noise level, at least 8000 seconds were simulated. For each noise level, the resulting bit distribution was fitted with a Gaussian-like distribution. This method enables accurate calculation of BERs as small as 10^{-30} , for which otherwise a huge number of simulations would be required.

Fig. 12 shows the resulting BER versus E_b/N_0 (1-second bit energy per spectral noise density):

- As expected, the matched filter detector performs best. It cannot be improved by additional filtering.
- The Goertzel and CIC detectors perform identically and require approx. 1.2 dB higher E_b/N_0 than the matched filter detector (for equal BER).
- The diode detector performs worse, even for narrow bandwidths. This is mainly due to the rectification which prevents noise averaging.
- The diode detector improves by reducing its input bandwidth (at the expense of precision).
- If a narrow input filter is not feasible, the diode detector can be improved by averaging.

It is concluded that the Goertzel and CIC detectors introduced in this paper provide nearly optimum performance in the presence of AWGN. This enables efficient DCF77 receivers with high immunity.

Fig. 13 compares some of the simulations from this paper with previous work [5].

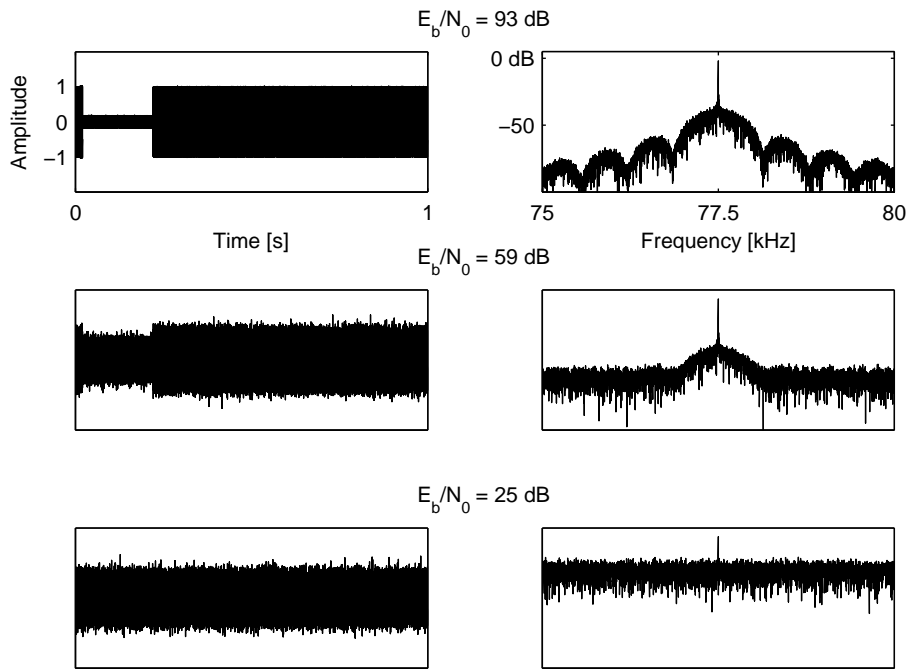


Fig. 11. Examples of DCF77 signal in the time- and frequency domain for different levels of noise. At the decoding limit of $E_b/N_0 = 2.7$ dB (CIC or Goertzel detector with maximum likelihood decoder, see Fig. 12), the signal is not visible anymore in a 1-second recording.

B. Narrowband interference

A realistic environment contains not only AWGN but also strong narrowband signals, for example electromagnetic interference from switched-mode power supplies or cathode ray tubes. A simulation was performed which uses a single sine-wave interference located at a frequency offset $\pm\Delta f$ from the carrier. The simulation uses a binary search algorithm to find the interference amplitude where the BER reaches a predefined value, in this case 0.34 (maximum of 1-hour ML-decoder, as will be seen later in Fig. 22).

Fig. 14 shows the resulting frequency responses:

- As expected, all detectors are more sensitive close to the carrier.
- The Goertzel detector performs better than expected from its seemingly poor selectivity (Fig. 7): Due to the PM pseudo-random correlation, interference at a large enough frequency offset phase-wraps and cancels.
- The CIC detector performs best, even better than the matched filter detector. The matched filter detector is optimal only against AWGN, not against narrowband interference.
- Inside the main lobe and first side lobe (up to ± 1292 Hz from the carrier), all detectors perform nearly identically. The CIC detector's better selectivity becomes useful only outside this frequency range.

It is concluded that a detector's narrowband interference rejection is not only dependent on its input filter, but also on its sensitivity on the signal spectrum.

C. Assumptions

The detector simulations described above are based on the following assumptions:

- The DCF77 signal is modelled with AM (including the $250 \mu\text{s}$ blanking interval) and PM. The transmit antenna response was reconstructed from time-domain plots in [1].
- Noise is modelled as AWGN, narrowband interference as a sine wave.
- The receiver has a perfect frequency reference.
- Second synchronisation was successful. This is a reasonable assumption because second synchronisation can be time-averaged for as long as required, whereas second detection can use data from only 1 second.
- Multipath propagation due to ground wave, sky wave, and re-radiation on conductive objects [8] is ignored.
- The receiver antenna is modelled for the Demonstration Receiver. For the other receivers, an antenna with linear phase is assumed.

V. TIME DECODER

The detectors described above convert the DCF77 radio signal into a stream of 1 bit/s (assuming identical AM and PM bits). This bit stream is read by the decoder which calculates the current time. This time is then used to update the receiver's local clock.

A. Assumptions

The analysis below assumes a clock which shows hour, minute, and seconds. For simplification, the following information provided by DCF77 is ignored:

- Date: The date is constant throughout 24 hours and therefore easy to decode.
- Daylight saving, leap second: The change between summer time and winter time, and the introduction of a leap

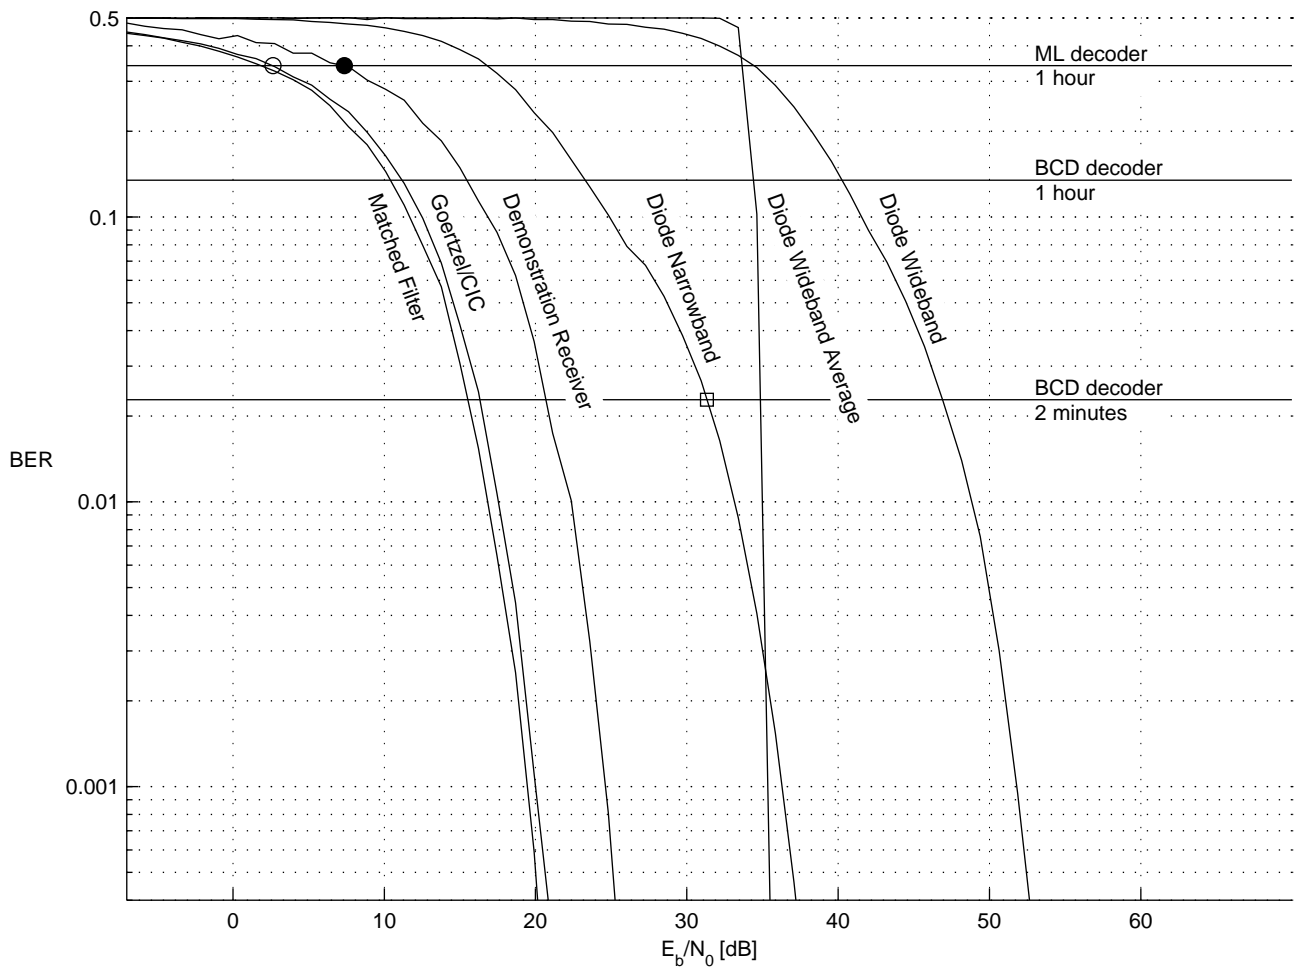


Fig. 12. Detector comparison for AWGN (simulation results). The markers indicate the maximum noise level where time reception is still possible: Non-filled circle at $E_b/N_0 = 2.7$ dB for the Goertzel or CIC detector with 1-hour ML-decoder. The filled circle at $E_b/N_0 = 7.4$ dB for the Demonstration Receiver. The square at $E_b/N_0 = 31.3$ dB for conventional low-cost receivers. The horizontal lines are the time decoder BER limits, assuming a decoding probability $p_{ok} = 0.5$ (see section V-B)




Fig. 13. Comparison of AWGN detector simulations with [5]. The dashed lines are simulations from this paper, the solid lines analytical derivations from [5]. The thick line is a PLL-receiver simulation from [5].

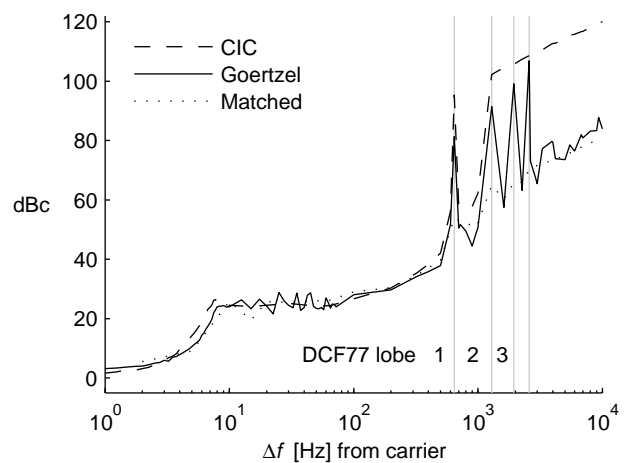


Fig. 14. Narrowband interference. The vertical axis shows the power (at the detector input) of a sine-wave interference located at a frequency offset $\pm\Delta f$ from the carrier, such that the BER reaches the decoding limit of the ML-decoder.

second are each announced with a DCF77 bit during 1 hour.

- Weather/alarm data: The weather data format is proprietary. It is expected that the format is optimised for net data throughput and contains only little redundancy. Therefore, decoding the weather data benefits from an improved detector, but not from an improved decoder.

B. Figures of merit

To compare different time decoders, the following figures of merit are relevant:

- p_{ok} is the probability of correct decoding, depending on the BER. For a noise-free signal and therefore a bitstream with $BER = 0$, any decoder should have $p_{ok} = 1$. For increasing BER, p_{ok} degrades, but not necessarily monotonically. In the following analysis, BER_{max} is defined as the smallest BER where p_{ok} reaches a predefined limit of 0.5.
- p_{off} is the probability of incorrect decoding, and also depends non-monotonically on the BER, with $p_{off} = 0$ at $BER = 0$. When designing a decoder, p_{off} must be minimised over all BER, even at the expense of BER_{max} , because decoding the wrong time is worse than not decoding.

The limit of $p_{ok} = 0.5$ means that 1 out of 2 reception attempts succeeds. This is arbitrary, and any other value could be used as well. For consumer receivers with a typical low-cost quartz oscillator and one reception attempt per day, $p_{ok} = 0.5$ is sufficient to remain accurate to within approx. 1 s. A receiver which requires higher accuracy should synchronise continuously (as implemented in the Demonstration Receiver).

The probability of detected error is $1 - p_{ok} - p_{off}$ and is not considered further.

C. BCD decoder (existing)

The BCD (binary coded decimal) decoder is a simple and probably the most frequently used algorithm to decode the DCF77 time. It can be implemented as follows:

- 1) Wait for the minute start.
- 2) Record 1 minute.
- 3) Verify parity bits.
- 4) Verify value of minute (0 to 59) and hour (0 to 23).
- 5) Repeat for 1 more minute.
- 6) Verify the minute sequence.
- 7) If any step fails, try again.

1) *Error detection:* The parity, range, and sequence checks are quite robust:

- All odd-numbered bit errors are detected by a parity check.
- 2-bit errors are detected by the parity check (if they occur in different parity sections) or by the sequence check.
- Only 4 and higher even-numbered bit errors in certain positions may go unnoticed.

An example of decoding with bit errors is shown in Fig. 15. On the first line, a single-bit error in minute 11:48 is detected by the parity check. Decoding succeeds two error-free minutes




Fig. 15. An example of decoding with bit errors.

later. On the second line, a two-bit error changes 11:48 into 11:00, which is not detected by the parity check, but by the sequence check:

2) *Analysis:* For the following analysis, it is assumed that second and minute synchronisation have succeeded. An alphabet of $|C| = 60 \times 24$ codewords (minutes per day), each consisting of $L = 30$ bits (bits 21–35 of 2 sequential minutes), is used. It follows:

$$p_{ok} = (1 - BER_{max})^L \stackrel{!}{=} 0.5 \quad (2)$$

$$BER_{max} = 0.023 \quad (3)$$

$$\max_{\forall BER} (p_{off}) = \frac{1}{|C|} \sum_{\substack{\forall x,y \in C \\ x \neq y}} BER^{D_{xy}} (1 - BER)^{L - D_{xy}} \quad (4)$$

$$= 1.8 \cdot 10^{-4} \quad \text{at } BER = 0.13 \quad (5)$$

where D_{xy} is the Hamming distance between codewords x and y , which is at least 4 (due to the parity bits of minute and hour). This calculation was performed with an exhaustive codeword search.

Due to the small number of codewords, the probability of undetected error given random data is:

$$p_{off}(BER = 0.5) = \frac{|C|}{2^L} = 1.3 \cdot 10^{-6}, \quad (6)$$

which is much smaller than the maximum p_{off} given above.

For a fair comparison of the BCD decoder with the maximum-likelihood decoder (described below), both must be run for the same maximum duration of 60 minutes. The BCD decoder simply tries again until at most 60 minutes have passed. Since analytical expressions become complicated, a simulation was used to obtain the following result:

$$BER_{max,60} = 0.13 \quad (7)$$

$$\max_{\forall BER} (p_{off,60}) = 8.4 \cdot 10^{-3} \quad \text{at } BER = 0.16 \quad (8)$$

$p_{off,60}$ is nearly 1%, which is unacceptably high. However, p_{off} can be easily improved by additional checks on the received signal, e.g. a parity and range check of the date bits.

D. Maximum-likelihood decoder (new)

The maximum-likelihood (ML) decoder is a new DCF77 time decoder algorithm introduced by this paper. The ML decoder can deal with much higher BER than the BCD decoder described above. The ML decoder was implemented on the Demonstration Receiver described in section VI.

The ML decoder is based on the following idea: The DCF77 signal is highly redundant. It may be interpreted as 1 bit per second or 1 codeword per minute, but due to the fixed and known sequence of time it really consists of only 1 very long

	⚡ ⚡	⚡	⚡ ⚡⚡	⚡ ⚡
Received:	1x:4x	11:x8	x1:xx	1x:5x
Most likely:	11:47	11:48	11:49	11:50

Fig. 16. An example sequence with bit errors which can still be decoded by the maximum-likelihood decoder.

codeword (400 years, the length of the Gregorian calendar). The goal of a receiver is therefore not to decode this single codeword, but to synchronise to it. A number of initially unknown bits is received, but as soon as the time is decoded, all future bits are known (ignoring leap seconds).

The ML decoder records the received signal and calculates the time that was most likely sent. The recorded data are never thrown away (only if memory overflows), and are always analysed as a whole. Since the codeword is not actually decoded, neither the BCD format nor the parity bits are of importance, and any other coding with equal or higher Hamming distance would work as well.

The number of bit errors which can be corrected is up to half the Hamming distance (although at this limit the probability of undetected error is high). See Fig. 16 for an example sequence with bit errors which can still be decoded:

1) *Algorithm:* An ideal decoder would find the maximum correlation between the received signal and the known sequence of 24 hours \times 60 minutes \times 60 seconds. If the date were also decoded, the known sequence would be 400 years long. A full correlation is computationally intensive, therefore a partial correlation is preferred which sequentially searches for the maximum correlation of second, minute, and hour:

- 1) Find the most likely minute start position by correlating with the known pattern of constant bits 0–14 (PM), 20, and 59 (Fig. 17). Also, all parity-checked sections can be added to the correlation sum by sequential multiplication, without actually decoding them. This results in the current second.
- 2) Based on the known minute start, find the most likely minute (Fig. 18).
- 3) Based on the known minute, find the most likely hour. Depending on the minute, an hour-crossing must be taken into account (Fig. 19).

The partial correlation is nearly as good as the full correlation. A simulation of both methods shows that for equal decoding probability p_{ok} , the BER degradation is less than 0.01 (additive). This can be explained as follows: The second correlation uses 17 bits, whereas the minute and hour correlations use only 8 and 7 bits. Therefore, it is unlikely that the maximum full correlation occurs in a position which is not also the maximum second correlation. The same argument is valid for the minute.

Given a sequence of bits, the ML decoder returns the most likely time, which is not necessarily the correct time. If the BER is too high, or there is no DCF77 signal at all, the decoder should indicate this and wait for more data. The ML decoder as implemented therefore includes a confidence check which analyses the correlation peaks of second, minute, and hour.

The probability of undetected error p_{off} is determined by the parameters of the confidence check, and can be chosen arbitrarily. A reasonable trade-off must be found between low p_{off} and high p_{ok} . The Demonstration Receiver is designed with a maximum $p_{off} = 5.5 \cdot 10^{-5}$, which is considered low enough for practical use.

2) *Soft bits:* At the DCF77 transmitter, each second is assigned either a 0 or 1 (neglecting the difference between AM and PM coding). At the receiver, the noisy signal is processed as real numbers representing amplitude or phase. At the input of the time decoder, it seems natural to apply a threshold to determine whether the received second is a “hard” bit 0 or 1.

If no threshold is applied, the value can be interpreted as a “soft” bit in the range -1 to $+1$, where -1 indicates a bit 0, and $+1$ indicates a bit 1. A value of 0 indicates an undecided bit which is equally likely a bit 0 or 1. See Fig. 20 for example soft bit distributions.

The use of soft bits is motivated by the effect of AWGN on the detector output. For all detector types discussed above, the soft bit distribution is Gaussian. For moderate levels of noise, bits which change sign due to noise remain mostly close to 0, indicating an undecided bit. If a threshold were applied, such a bit would be indistinguishable from a flipped bit.

Using soft bits instead of hard bits improves time decoding performance. For equal decoding probability and equal recording duration, simulations show that the soft bit decoder can handle approx. 0.066 higher BER (additive). The hard-bit decoder is less robust against undetected errors and thus requires a more stringent confidence check, which further degrades its performance.

Soft bits require more memory than hard bits, but the complexity of the time decoder algorithm remains the same.

3) *Partial minute decoding:* When a DCF77 receiver is powered up, it does not yet know the current second. The BCD decoder described in the previous section first waits for the minute start and only then starts recording bits. Decoding takes between 2 to 3 minutes.

In contrast, the ML decoder starts recording bits as soon as it is powered up. To decode the time, the minute start (AM bit 59 or PM bits 0–9), minute (21–28), and hour (29–35) are required. These bits need not necessarily be from the same minute, as shown in Fig. 21.

With a good signal, decoding the time takes as little as 60 s for any powerup time. In the best case, if the receiver is powered up at second 0, decoding takes only 35 s. With a bad signal, the ML decoder waits until the recorded data fulfil the confidence checks.

Partial minute decoding comes “for free” with the ML decoder and requires no additional algorithm.

[9] describes the same functionality, but without details of the underlying algorithm.

4) *Analytical approximation:* This section calculates the performance of the ML decoder, based on the following simplifying assumptions:

- Second and minute synchronisation have succeeded.
- The same minute is repeated m times (instead of consecutive minutes).
- Received bits are either 0 or 1 (no soft bits).

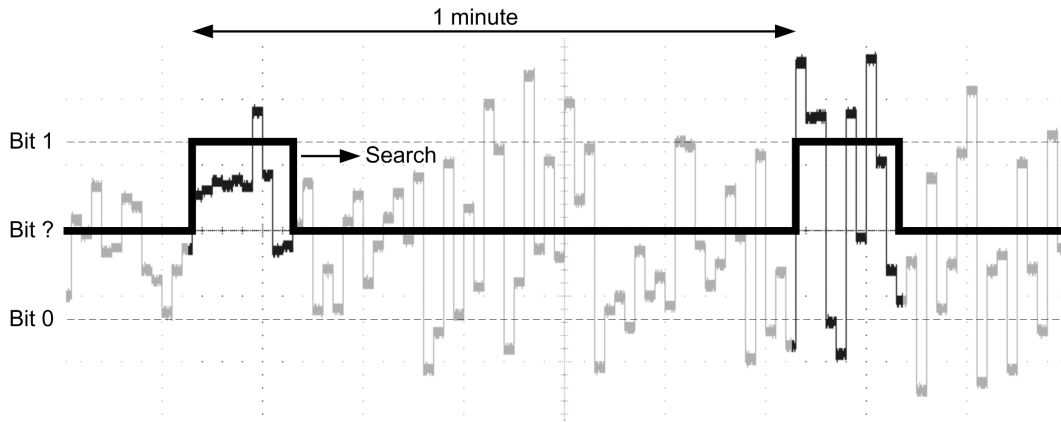


Fig. 17. Maximum-likelihood decoding of second (simplified): The minute start pattern of $10\times$ PM bit 1 is correlated with the received signal. The position of the correlation maximum is equivalent to the current second.

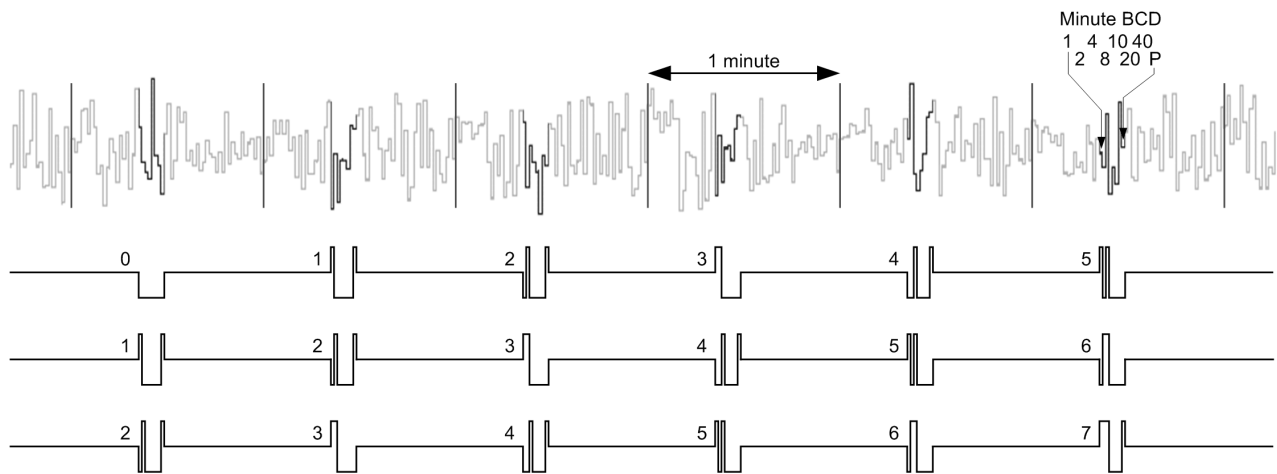


Fig. 18. ML decoding of minute: Once the second is known, the minute is correlated. The first line calculates the correlation of the received signal with minute 0, the second line with minute 1, etc. The correlation uses only bits 21–28 (minute BCD bits 1 through 40 and parity bit). The parity bit is treated no differently from the BCD bits.

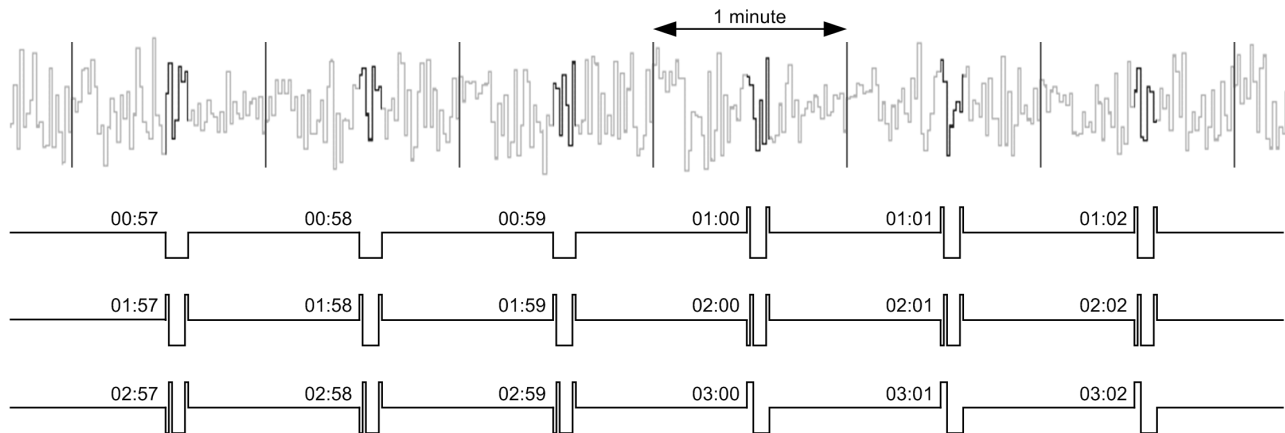


Fig. 19. ML decoding of hour: Once the minute is known, the hour is correlated. Depending on the minute, an hour crossing must be taken into account (as shown).

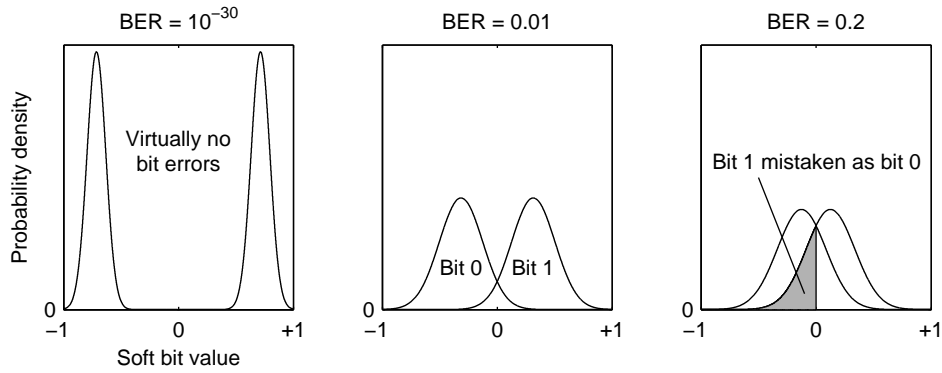


Fig. 20. Example soft bit distributions for different bit error rates.




Fig. 21. Partial minute decoding: The ML decoder is capable of combining minute segments. This enables time decoding within only 60 s after powerup.

There are $24 \times 60 = 1440$ distinct minutes, which could be represented by 10.5 bits. The DCF77 BCD code uses 15 bits (including 2 parity bits). Subsequent minutes use a different representation for the same information, which reduces the inefficiency of the BCD code somewhat. Therefore, by rounding up the 10.5 bits, a $(n = 15, k = 11)$ code is used for Hamming distance calculations.

Ignoring the parity bits, repeating a minute m times may also be regarded as repeating each of the n bits individually. For such an m -repetition code of a single bit, up to $\lfloor m/2 \rfloor$ bits may be wrong, resulting in the following probability of error per bit:

$$p_1 = 1 - \sum_{i=0}^{\frac{m-1}{2}} \binom{m}{i} \text{BER}^i (1 - \text{BER})^{m-i} \quad (9)$$

Taking into account the $n - k$ parity bits, the probability of correct decoding is:

$$p_{\text{ok,ML}} = \sum_{i=0}^{\frac{n-k}{2}} \binom{n}{i} p_1^i (1 - p_1)^{n-i} \quad (10)$$

5) *Sequential ML decoder*: A similar algorithm is described in [10]:

- 1) 60 accumulators correlate each bit of a minute with the minute start bit. The largest accumulator corresponds to the minute start and thus the current second.
- 2) Once the second is found, 2 accumulators correlate the toggling of the least significant minute bit.

3) This continues with the next bit, etc.

Advantages compared with ML decoder:

- The correlation length is independent of memory size. In principle, an infinitely long correlation is possible, which should be able to cope with higher BER (though [10] contains no performance analysis).

Disadvantages:

- Longer time until first synchronisation.
- Loss of synchronisation causes loss of history: If an accumulator maximum moves due to noise, all following accumulators are reset.

E. Performance comparison

Fig. 22 shows the performance of the ML and BCD time decoders. This figure was obtained through simulations: For a given BER and number of minutes, a noisy minute sequence is generated and fed to a decoder model which calculates the time, either correctly or not. Multiple runs approximate the decoding probability p_{ok} and p_{off} . This is repeated for all BERs and numbers of minutes, resulting in a surface whose crossing at $p_{\text{ok}} = 0.5$ provides the lines of Fig. 22.

The analytical model $p_{\text{ok,ML}}$ agrees well with the simulation model of the ML decoder. Introducing confidence checks degrades BER_{max} but reduces the chances of incorrect decoding.

It is concluded that the new ML decoder introduced by this paper outperforms the existing BCD decoder by far. At the maximum $\text{BER} = 0.34$, the bit sequence is nearly indistinguishable from noise, but after 1 hour, the ML decoder successfully displays the time.

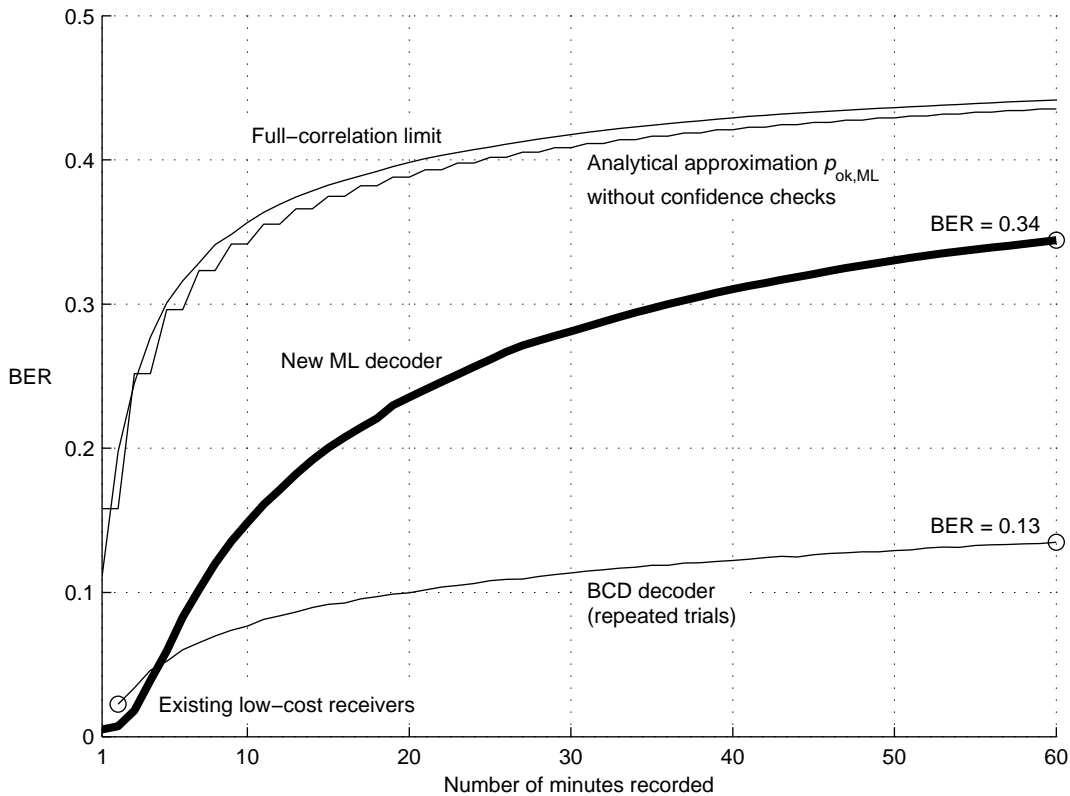


Fig. 22. Time decoder performance (simulation results for decoding probability $p_{ok} = 0.5$). The ML decoder introduced by this paper performs nearly as good as the theoretical optimum full-correlation. By introducing confidence checks, the probability of incorrect decoding is reduced significantly. The ML decoder was implemented on the Demonstration Receiver described in section VI.

VI. DEMONSTRATION RECEIVER

To confirm the capabilities of the algorithms introduced in this paper, the Goertzel-PM detector and the ML-decoder with 1-hour history were implemented on an FPGA-based Demonstration Receiver. Fig. 23 and 24 show the custom-built PCB (printed circuit board) and its block diagram.

The 9 cm wide PCB has an off-the-shelf ferrite rod antenna mounted on a pole. An initial JFET amplifier senses the high-impedance antenna signal, the following bandpass filter reduces interference. A programmable gain amplifier (controlled by the FPGA) automatically adjusts the gain. The ADC (analog-to-digital converter) samples at $12f_c$. The FPGA reads the ADC data, then processes the Goertzel-PM detector and ML-decoder on the signal. Finally, the time is shown on a display.

From simulations of the receiver, including antenna, analog input stage, clock correction, synchronisation, detector, and decoder, it is expected that the receiver works with $E_b/N_0 > 7.4$ dB (Fig. 12) and $BER_{max} = 0.34$ (Fig. 22).




Fig. 23. Demonstration Receiver with FPGA implementing the Goertzel-PM detector and ML decoder. This custom PCB was developed to verify the algorithms introduced by this paper.

A. Clock correction

The receiver clock frequency is usually less accurate than the atomic clocks which maintain the DCF77 time, with the following consequences:

- If the receiver clock runs too fast, its time may not be monotonic, which is undesirable. For example, it may

show 12:00:00, 12:00:01, then synchronise and show 12:00:00 again.

- A receiver rejects noise mainly by averaging the received signal for as long as possible. The receiver clock error shifts the averaging window which limits the averaging duration and thus performance.

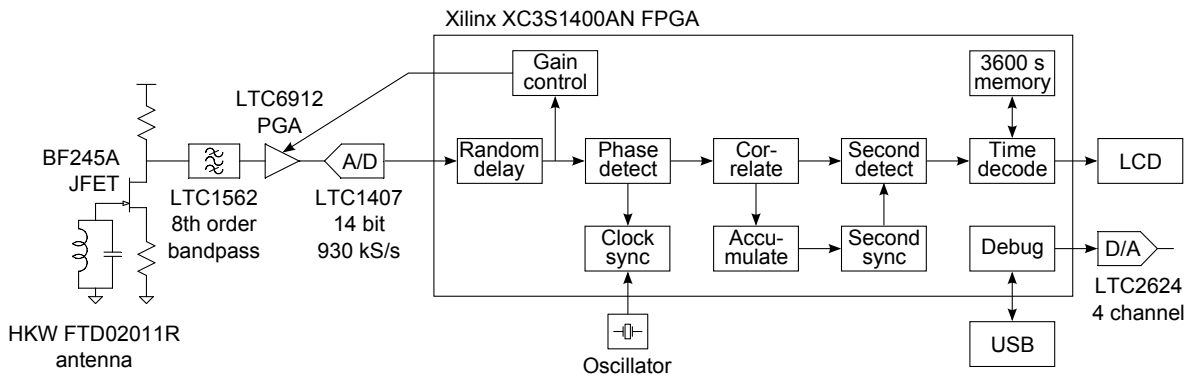


Fig. 24. Block diagram of Demonstration Receiver implementing the Goertzel-PM detector and ML decoder.

For the Goertzel-PM detector, simulations show that with 1 ppm clock error, 0.5 dB higher E_b/N_0 is required for equal BER. At 10 ppm, this raises to 6 dB. Uncompensated quartz oscillators usually have a tolerance of ± 20 ppm to ± 50 ppm. Therefore, for good performance a more accurate clock is required, such as:

- Compensated oscillator (e.g. TCXO or OCXO) which is inherently accurate.
- Low-cost oscillator synchronised to the accurate DCF77 carrier frequency, either analog (VCXO) or digital (as described below).

Fig. 25 and 26 show the clock correction algorithm as implemented on the Demonstration Receiver: A fast clock is normally divided by $d = 8$, which is adjusted to $d \pm 1$ for 1 out of N cycles, depending on the current correction factor. This enables a resolution of 0.003 ppm. For the Demonstration Receiver, simulations show that a clock accuracy of 0.1 ppm is sufficient for second synchronisation with $E_b/N_0 > 0$ dB.

Clock correction introduces distortion and ADC clock jitter: The SNR (signal-to-noise ratio) is limited to 50 dB, and the PM timing is changed by up to 0.1%. Both effects are negligible.

The phase comparison uses the Goertzel algorithm with a variable scaling constant as described in section II-B. After each step of the clock correction binary search, the Goertzel scaling constant is adjusted to narrow the input bandwidth of the clock correction. With increasing clock accuracy, noise immunity is improved. This enables longer averaging and thus even higher accuracy. If the DCF77 signal quality is slowly changing (for example with day and night), the clock correction locks when the signal is good and due to the higher immunity remains locked when the signal quality degrades.

A moving receiver incurs a Doppler frequency shift of $f_0(1 + \frac{v_p}{c})$, where v_p is the velocity component in the direction of the transmitter. A car driving at typical highway speeds causes a frequency shift of approx. 0.1 ppm, which is equivalent to the precision of the clock correction algorithm and is thus easily compensated.

A similar work was performed by [11], which measures the achievable frequency stability by synchronising to the DCF77 carrier and qualitatively observes the impact of interference and atmospheric noise.

B. Noise analysis

The received signal is degraded by external and internal noise. Internal noise is generated by the receiver itself. External noise can be classified into atmospheric, galactic, and man-made [12].

Galactic noise is relevant only at 3 MHz and higher. Man-made noise is highly dependent on location and is therefore modelled together with other noise.

1) *Atmospheric noise:* Atmospheric noise is caused by worldwide thunderstorms and is dependent on location, season, and time of day. For Europe, an all-year maximum of approx. $9 \text{ dB}(\mu\text{V}/\text{m}/\sqrt{\text{Hz}})$ is assumed (extrapolated from [8], [12]). At noon (all seasons), at least 16 dB less can be expected.

Atmospheric noise is only partially Gaussian. Lightning discharges cause short pulses which should be handled in the time-domain. One method is known as “hole-punching” which is used by the clock-correction algorithm. The detector does not require such processing, since only a single second bit is affected which is safely handled by the ML time decoder. The remaining part of the atmospheric noise is treated as Gaussian with the spectral density given above.

2) *Receiver-internal noise:* A Spice [14] noise simulation was performed including the voltage regulators, antenna, and analog input stage. To compare the internal noise with the DCF77 signal and the atmospheric noise, the internal noise at the ADC input is converted into an equivalent electric field strength, resulting in $6 \text{ dB}(\mu\text{V}/\text{m}/\sqrt{\text{Hz}})$ at the carrier (for the maximum programmable gain where noise is highest).

Away from the carrier, the noise decreases due to the antenna resonance and the bandpass filter. The equivalent field strength is therefore frequency-dependent, but for simplicity the pessimistic assumption is made that the maximum noise is constant for all frequencies.

The noise simulation was verified with a spectrum analyser measurement for all gain settings. The readings were offset by -5.0% , probably due to limited measurement bandwidth and pessimistic simulation assumptions. With this offset corrected, the readings are within $\pm 1.2\%$ of the simulation. Further measurements of total RMS noise showed that only 3 out of 6 randomly selected digital multimeters provide “True RMS” capability for such narrowband noise.




Fig. 25. Clock correction principle. The clock correction algorithm enables the Demonstration Receiver to synchronise its low-cost ± 50 ppm quartz clock to within 0.1 ppm of the DCF77 reference frequency. Such an accurate clock is needed for long-term signal averaging and thus high noise suppression.

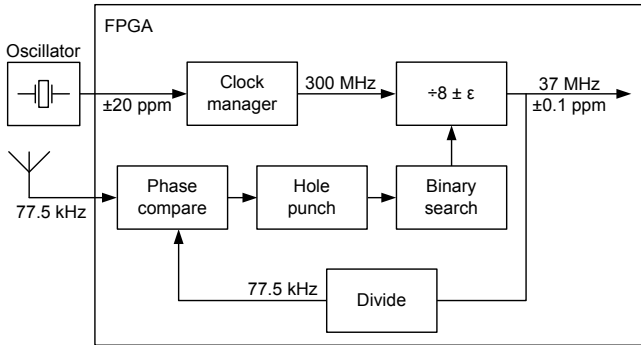


Fig. 26. Clock correction implementation: The “hole punch” algorithm mutes excessively large phase deviations caused by burst noise. The binary search uses only small steps such that a wrong decision (due to noise) can be corrected later.

C. Distance from transmitter

The allowable distance of a receiver from the DCF77 transmitter is determined by the receiver gain, noise, ADC resolution, and detector performance.

At the minimum distance, the strong DCF77 signal should not overload the receiver. A maximum signal strength of 100 dB(μ V/m) plus 20 dB of headroom is assumed, which determines the minimum gain of the receiver.

At the maximum distance:

- The receiver gain must amplify the antenna signal for the ADC. Simulations show that a signal amplitude of 1 LSB (least-significant bit) is sufficient, even with high noise levels. 2LSB was chosen for the design. It may be possible to increase the effective ADC resolution using dither, either through radio noise or added artificially, but this has not yet been investigated.
- A quiet location is assumed with only atmospheric (and internal) noise. Any additional external noise would decrease the range. With the noise level and the receiver’s E_b/N_0 limit known, the minimum DCF77 field strength can be calculated.

If the maximum distance of a receiver is to be determined experimentally, by actually travelling with the receiver away from the transmitter, the separate propagation of ground wave and sky wave must be considered. As the arrival delay between the two may be only a few carrier cycles [1], a high-resolution receiver and a quiet reception site are required.

Table I shows the resulting distances for different receiver types. The distance calculation considers only the ground wave. Including the sky wave, even a simple diode receiver

should work at up to 2000 km once per day [1], though with reduced accuracy.

It is concluded that atmospheric noise in the DCF77 band is relatively low, such that even simple receivers can cover Western and Central Europe. With improved algorithms as described in this paper, the range can be extended to cover the entire continent. The internal noise of the Demonstration Receiver is below the atmospheric noise level and degrades the reception range only slightly.

D. Accuracy and precision

The accuracy and precision of a DCF77 receiver depend on the following factors:

- Radio wave propagation: If the distance d to the transmitter is known, the propagation delay can be roughly compensated by d/c . For higher accuracy, longwave propagation models need to be considered [1], [8]:
 - Ground wave propagation depends mainly on ground conductivity and terrain. The exact propagation delay at a given location must be determined by measurement (e.g. using GPS). The delay variation is small and can be neglected compared to the achievable precision of DCF77 receivers.
 - Sky wave propagation becomes dominant at distances above approx. 1000 km. The propagation delay varies with the number of ionosphere hops, time of day, season, weather, and sun spot activity, and can change within only minutes. The sky wave enables longer range, but with reduced accuracy.
 - Re-radiation due to conductive objects causes multipath propagation [8], but the introduced delay is below the DCF77 precision and therefore negligible.
- The L-C resonance frequency of a typical tuned ferrite rod antenna has a production tolerance of approx. ± 200 Hz, which results in a group delay variation of approx. 100μ s. The delay can be factory-calibrated, but changes with temperature and age. This can be avoided with a wideband loop antenna [4] or an E-field whip antenna [5], both without resonance.
- Narrow input filters improve interference rejection, but limit the achievable precision of the AM falling edge [4] and the PM correlation.
- Discrete-time processing must trade off precision with the cost of processing and memory. The precision of the Demonstration Receiver is $3875 \text{ Hz} = 1/258 \mu$ s. This could be increased until approx. $1/f_c = 13 \mu$ s, a single carrier cycle [1].

Detector	Sync	Decoder	Performance E_b/N_0 [dB]	Minimum field strength [$\mu\text{V}/\text{m}$]	Distance [km]
Diode Narrowband	AM falling-edge	BCD 2-minute	31	180	1200
Goertzel-PM	Single-second	ML 1-hour	7.4	12	2200
CIC/Goertzel AM+PM	Full-minute	ML 1-hour	2.7	7	2400

TABLE I. Maximum distance from transmitter due to atmospheric and internal noise, assuming ground wave propagation. The underlying simulations use the antenna and analog input stage of the Demonstration Receiver, therefore these results provide realistic estimates. A receiver’s performance (E_b/N_0) determines how small the DCF77 signal can be compared to the noise, and therefore directly affects the distance.

- Clock frequency error: A receiver with a free-running clock incurs a time offset proportional to its frequency error and the synchronisation interval. Even a small frequency error of 0.1 ppm accumulates to a time offset of 9 ms per day. To minimise this time offset, the Demonstration Receiver synchronises every second.

In total, the Demonstration Receiver accuracy is designed to be approx. 1 ms, less than that after calibration.

The second start signal of the Demonstration Receiver was compared with a GPS receiver, by performing 54 runs from powerup to first DCF77 synchronisation. The readings fell within 362 μs , which is 1.4 times worse than the designed accuracy. All readings were taken with the initial yet uncorrected clock error, and should improve after clock synchronisation.

E. Clock leakage

The following clocks are used on the Demonstration Receiver:

- The ADC samples at $12f_c$
- The Goertzel algorithm calculates results at f_c
- The PM is correlated at $f_c/20$
- The time decoder works at 1 Hz

These frequencies and their harmonics leak to the antenna (e.g. through the power supply cable), where they are picked up and amplified. This masks the desired signal and degrades reception.

To solve this problem, the FPGA buffers ADC samples and processes them in bursts of random length. This spreads and shifts the clock leakage spectrum away from f_c , enabling undisturbed reception.

F. Observations

The following observations were made in an office building in a room with PCs, laboratory equipment, and fluorescent lighting: 4 different low-cost receivers (weather station, alarm clock, wristwatches) and a self-built diode detector receive the DCF77 time only at 0.5 m from a window, whereas the Demonstration Receiver synchronises correctly at approx. 7 m from the window.

The DCF77 signal quality degrades quickly with increasing distance from the window. This can be shown with the Demonstration Receiver, e.g. with a spectrum analyser after the antenna amplifier, or on one of the FPGA-internal signals from the clock correction or phase correlation algorithms.

At a quiet reception site, the DCF77 quality can be easily measured with a spectrum analyser. At a noisy site, this is

more difficult: The DCF77 signal’s high redundancy makes it detectable under so much noise that a simple visual inspection in the time or frequency domain is not useful (Fig. 11). Instead, the BER can be estimated, either from the soft bit distribution [13] or by comparing the received bits with the decoded bits (which only works after reception has succeeded). Knowing the BER and the detector model, E_b/N_0 can be computed.

VII. FUTURE WORK

A. Indoor radio environment

The performance analysis in this paper is based mainly on AWGN. The next step is to measure and model a realistic indoor radio environment, characterised by:

- Signal attenuation depending on location and building construction.
- EMI (electromagnetic interference) of switched-mode power supplies and other devices. This is governed by EMI regulations, which allows an approximation of the receiver’s minimum distance from a source of EMI.

B. New “Instant” DCF77 modulation

Indoor noise is expected to be dynamic, e.g. as overhead lights are switched on and off, or as computer load and with it power consumption varies. Therefore, instead of trying to deal with noise (as described in this paper), it may be better to avoid it by more frequent reception attempts. To increase the chances of success and to keep power consumption low, a reception attempt should be as short as possible.

Therefore, a new DCF77 modulation could be investigated which enables time reception in only a few seconds (instead of minutes). The existing AM and PM should not be disturbed, so an additional differential phase modulation on top of the pseudo-random code could be introduced. In this way, one DCF77 second would contain the full time information.

ACKNOWLEDGEMENTS

This paper is a private work. Many thanks to Dr. M. Weisenhorn (ZHAW) for helpful discussions, R. Wacker (Zühlke Engineering AG) for laboratory equipment and funding of the Demonstration Receiver, Dr. A. Bauch (PTB), Dr. D. Piester (PTB), P. Kamp (FreeBSD), and Dr. P. de Boer (UTwente) for their support on DCF77.

Bit	AM	PM
0	0	1
1	Weather/alarm	1
2		1
3		1
4		1
5		1
6		1
7		1
8		1
9		1
10		0
11		0
12		0
13		0
14		0
15	Call/Antenna	
16	A1 (Change of Z1, Z2)	
17	Z1 (0 = CET, 1 = CEST)	
18	Z2 = $\overline{Z1}$	
19	A2 (Leap second)	
20	1	
21	Minute	1
22	Minute	2
23	Minute	4
24	Minute	8
25	Minute	10
26	Minute	20
27	Minute	40
28	P1 = xor(21:27)	
29	Hour	1

Bit	AM	PM
30	Hour	2
31	Hour	4
32	Hour	8
33	Hour	10
34	Hour	20
35	P2 = xor(29:34)	
36	Day	1
37	Day	2
38	Day	4
39	Day	8
40	Day	10
41	Day	20
42	Weekday	1
43	Weekday	2
44	Weekday	4
45	Month	1
46	Month	2
47	Month	4
48	Month	8
49	Month	10
50	Year	1
51	Year	2
52	Year	4
53	Year	8
54	Year	10
55	Year	20
56	Year	40
57	Year	80
58	P3 = xor(36:57)	
59	(no AM)	0

TABLE II. DCF77 AM and PM encoding.

REFERENCES

[1] D. Piester et al., “Dissemination of time and standard frequency using DCF77” (original title: “Zeit- und Normalfrequenzverbreitung mit DCF77”), PTB-Mitteilungen vol. 114 (Heft 4), pp. 345–368, 2004

[2] A. Bauch et al., “Dissemination of time and frequency using DCF77: 1959–2009 and beyond” (original title: “Zeit- und Frequenzverbreitung mit DCF77: 1959–2009 und darüber hinaus”), PTB-Mitteilungen vol. 119 (Heft 3), pp. 217–240, 2009

[3] P. Hetzel, “Time dissemination via the LF transmitter DCF77 using a pseudo-random phase-shift keying of the carrier”, Proceedings of the 2nd European Frequency and Time Forum (EFTF), Neuchâtel University, Neuchâtel, Switzerland, pp. 351–364, 16–18 March 1988

[4] P. Hetzel, “Longwave time dissemination using amplitude modulated time signals and pseudo-random carrier phase shifting” (original title: “Zeitübertragung auf Langwelle durch amplitudenmodulierte Zeitsignale und pseudozufällige Umtastung der Trägerphase”), Ph.D. dissertation, Dept. Prod. Eng., Univ. Stuttgart, Stuttgart, Germany, 1988

[5] J. Wietzke, “Criteria for the uniform classification of theoretical performance of time signal receivers, considering new digital variants” (original title: “Kriterien zur einheitlichen Beurteilung der prinzipiellen Leistungsfähigkeit verschiedener Zeitzeichenempfänger unter Einbeziehung neuer digitaler Varianten”), Ph.D. dissertation, Dept. Inf. Tech., Tech. Univ. Darmstadt, Darmstadt, Germany, 1988

[6] E. Jacobsen, R. Lyons, “The Sliding DFT”, IEEE Signal Processing Magazine, vol. 20, no. 2, pp. 74–80, Mar. 2003

[7] C. Kandziora, R. Weigel, “Low-Power Sub-Microsecond Time Synchronization”, private communication (not yet published)

[8] W. J. Pelgrum, “New Potential of Low-Frequency Radionavigation in the 21st Century”, Ph.D. dissertation, Faculty for Electrical Engineering, Techn. Univ. Delft, Delft, The Netherlands, 2006

[9] R. Mohr, M. Schubert, “Technology and development of radio-controlled clocks” (original title: “Funkuhrtechnik und Funkuhrentwicklung”), Journal “Wechselwirkungen”, Yearbook 2000, pp. 76–86, Univ. Stuttgart, Stuttgart, Germany, 2000

[10] M. Wierich, “A digital DCF77 receiver with high sensitivity” (original title: “Ein digitaler DCF77-Empfänger mit hoher Empfindlichkeit”), Diploma thesis, Dept. of Computer Science, Univ. Hamburg, Hamburg, Germany, 1998

[11] K. Kalliomäki et al., “Feasibility of DCF77 or NTP-Time Servers to Control Carrier Frequencies of Base Stations”, Frequency Control Symposium, 2007 Joint with the 21st European Frequency and Time Forum, 2007, pp. 865–867

[12] “Radio noise”, Recommendation ITU-R P.372-10, Oct. 2009

[13] L. T. Smit et al., “Soft Output Bit Error Rate Estimation for WCDMA”, Proceedings of Personal Wireless Communications Conference, Venice, Italy, Sep. 2003, pp. 448–457

[14] LTspice IV, Linear Technology Corporation



Daniel Engeler received the MSc in electrical engineering and information technology in 2004 from the Swiss Federal Institute of Technology (ETH), Zürich, Switzerland.

He is currently working as a hardware engineer at Zühlke Engineering AG designing industrial electronics.