Philips tech. Rev. 40, 174-179, 1982, No. 6

Digital-to-analog conversion in playing a Compact Disc

D. Goedhart, R. J. van de Plassche and E. F. Stikvoort

Introduction

The last stage in the series of operations on the signal in the Compact Disc system is the return from the digital code to the analog signal, which has the same shape as the acoustic vibration that was picked up by the microphone.

After decoding and error correction the digital signal has the form of a series of 16 bit words. Each word represents the instantaneous numerical value of the measured sound pressure in binary form, and is therefore a sample of the acoustic signal. There are 44 100 of these words per second.

The digital-to-analog converter in the Compact Disc player generates an electric current of the appropriate magnitude for each word and keeps it constant until the next word arrives. The electric current thus describes a 'staircase' curve that approximates to the shape of the analog signal (*fig. 1a*). In terms of frequency, the steps in the staircase represent high frequencies, which extend beyond the band of the analog audio signal (20 Hz - 20 kHz). These high frequencies have to be suppressed by a lowpass filter; in the Compact Disc player their level should be reduced to at least 50 dB below that of the maximum audio signal.

If this high attenuation of the frequencies above the audio band is to be achieved solely with an analog lowpass filter, the filter must meet a very tight specification. It was decided to avoid this problem in the Philips Compact Disc player by introducing a filter operation, earlier in the digital stages. This was done by 'oversampling' by a factor of four: a digital filter, operating at four times the sampling rate (4×44.1 kHz = 176.4 kHz) delivers signal values at this increased frequency, thus refining the staircase curve (fig. 1*b*) and making it easier to filter out the high frequencies. As a result it is possible to make do with a relatively simple lowpass filter of the third order after the digital-to-analog conversion.

The conversion of the 16 bit words into an analog signal is performed in the Philips Compact Disc player by a 14 bit digital-to-analog converter available as an integrated circuit and capable of operating at the high sampling rate of 176.4 kHz. Partly because of the fourfold oversampling and partly because of the feedback of the rounding-off errors in antiphase, rounding off to 14 bits does not result in a higher noise contribution in the audio band. This remains at the magnitude corresponding to a 16 bit quantization (signal-to-noise ratio about 96 dB), so that even though there is a 14 bit digital-to-analog converter it is still possible to think in terms of a 16 bit conversion system.

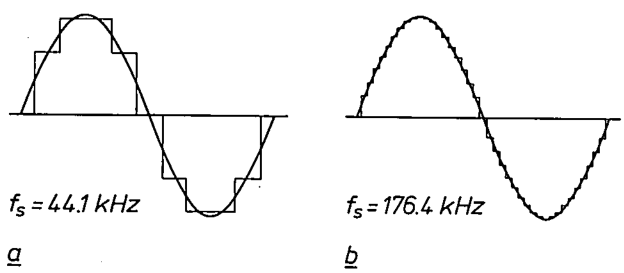

Fig. 1. A sinusoidal signal at 4.41 kHz sampled with a sampling rate f_s of 44.1 kHz (*a*) and with a frequency four times higher (*b*). In (*b*) the 'staircase' curve approximates more closely to the analog waveform, and the high frequencies present in the staircase signal are more easily filtered out.

In comparison with direct 16 bit digital-to-analog conversion, which must be followed by a lowpass filter with a sharp cut-off to give sufficient suppression of signals at frequencies above 20 kHz, our conversion system has a number of advantages. The first is the linear phase characteristic, which can be obtained with a digital filter, but not with an analog filter; the second is a filter characteristic that varies with the clock rate

Ing. D. Goedhart is with the Philips Audio Division, Eindhoven; Ir R. J. van de Plassche and Ir E. F. Stikvoort are with Philips Research Laboratories, Eindhoven.

and is therefore insensitive to variation in the speed of rotation of the disc. Finally, because the quantization steps are smaller, the maximum 'slew rate' that these circuits must be able to process is lower (the slew rate is the rate of variation of output voltage). There is therefore less chance of intermodulation distortion because the permitted slew rate has been exceeded.

44.1 kHz. The frequency spectrum of such a series is illustrated in fig. 3*b* [1]. In theory it is infinite; above the baseband 0-20 kHz can be seen integral multiples of the sampling frequency with their left-hand and right-hand sidebands. Between these bands there are transition regions, the first for example being between 20 kHz and 24.1 kHz.

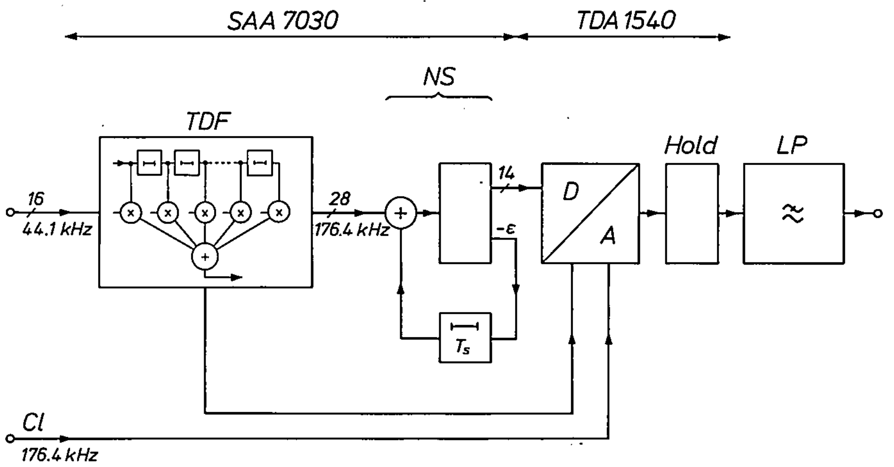

Fig. 2. Block diagram of the digital-to-analog conversion. *TDF* digital transversal filter which brings the sampling rate of 44.1 kHz to 176.4 kHz and attenuates signals in the bands around 44.1 kHz, 88.2 kHz and 132.3 kHz. *NS* noise shaper in which the rounding-off error is delayed by one period T_s after rounding-off to 14 bits and then fed back in the opposite sense. *D/A* 14 bit digital-to-analog converter. *Hold* hold circuit. *Cl* clock signal. *LP* lowpass 3rd-order Bessel filter.

The entire series of operations in the digital-to-analog conversion is shown as a block diagram in *fig. 2*. The oversampling takes place in the digital filter *TDF* to which the input signal is fed. The filter output signal is then rounded off to 14 bits, and the rounding error is fed back in the opposite sense in the noise shaper *NS*. The digital filter and noise shaper are located in a single integrated circuit in NMOS technology (type SAA 7030). This IC processes both stereo channels. Then follow the digital-to-analog converter *D/A* and a hold circuit, combined in a single IC type (TDA 1540) in bipolar technology; for each stereo channel there is a separate IC. The analog signal finally passes through a lowpass filter.

Suppression of frequencies above the audio band

Direct digital-to-analog conversion of the presented signal provides a series of analog signal samples (*fig. 3a*). These have the form of pulses that — in theory — are infinitely short, but have a content (duration times amplitude) corresponding to the sampled signal value. The repetition frequency is

This entire spectrum must not be passed on to the player amplifier and loudspeaker. Even though the frequencies above 20 kHz are inaudible, they would overload the player amplifier and set up intermodulation products with the baseband frequencies or possibly with the high-frequency bias current of a tape recorder. Therefore all signals at frequencies above the baseband should be attenuated by at least 50 dB.

To produce such an attenuation, an analog filter after the digital-to-analog converter will inevitably have to contain a large number of elements and require trimming. In addition a linear phase characteristic is required in the passband so that the waveform of pulsed sound effects will not be impaired. In the Philips Compact Disc player these requirements are met in a different way, by means of:

— fourfold oversampling of the signal in the digital phase,

— a digital filter operation,

[1] The principles of the digital processing of audio signals are explained in a very readable account by B. A. Blesser in Digitization of audio: a comprehensive examination of theory, implementation, and current practice, J. Audio Engng Soc. **26**, 739-771, 1978.

— a hold function after the digital-to-analog conversion,
— a third-order Bessel filter in the analog-signal path.

A digital transversal filter is used for the filtering after oversampling. To understand the operation of the filter, we can think of it as consisting of 96 elements (*fig. 4a*), while the delay in each element is $(176.4\times10^3)^{-1}$ s, i.e. a quarter of the sampling period or $\frac{1}{4}T_s$. Four times in each period the filter takes up new data. At three of these four times the content of this data is zero, since the oversampling is done by the introduction of intermediate samples of value zero. This means that only 24 of the 96 elements are filled at any one time. The contents of each element are multiplied by a coefficient c. The filter provides data at a rate of 176.4 kHz; each number is the sum of 24 nonzero multiplications. In this way the filter always calculates three new sample values at the locations of the zero samples.

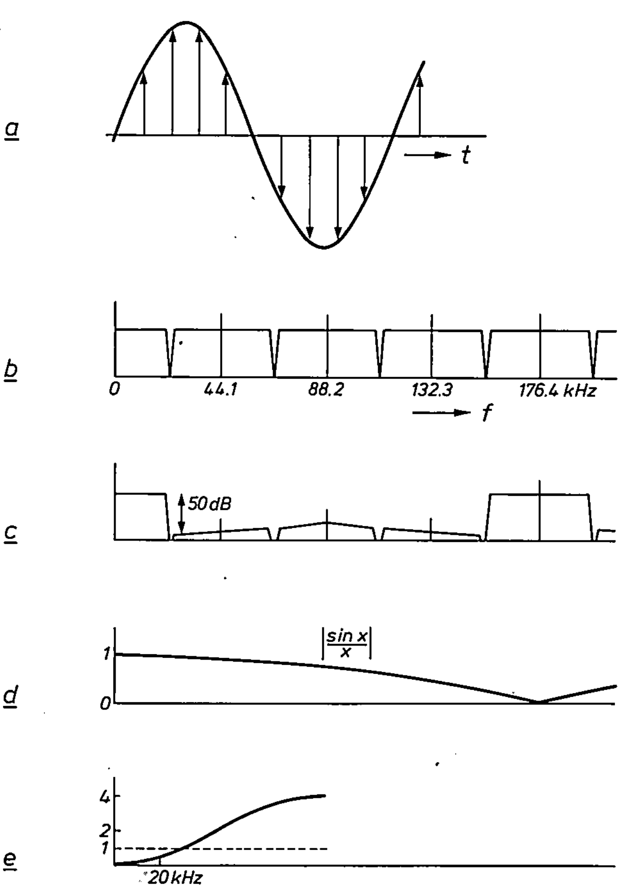

Fig. 3. *a*) A train of periodic pulses that sample an analog signal waveform. *b*) Frequency spectrum of such a pulse train. The pulse repetition frequency is 44.1 kHz, the sampled signal occupies the audio frequency band (0-20 kHz). *c*) Frequency spectrum for oversampling and filtering of the same signal at 176.4 kHz. It is now much easier to filter out the frequencies above the audio band. *d*) A hold circuit after the digital-to-analog converter keeps a signal sample at the same value until the arrival of the next sample. The frequency spectrum in *c* is thus multiplied by the function $|(\sin x)/x|$ with a first zero at 176.4 kHz. *e*) Noise spectrum after the noise shaper. In the audio range of interest the noise is considerably attenuated compared with the flat noise spectrum (*dashed line*) that would be obtained without noise shaping.

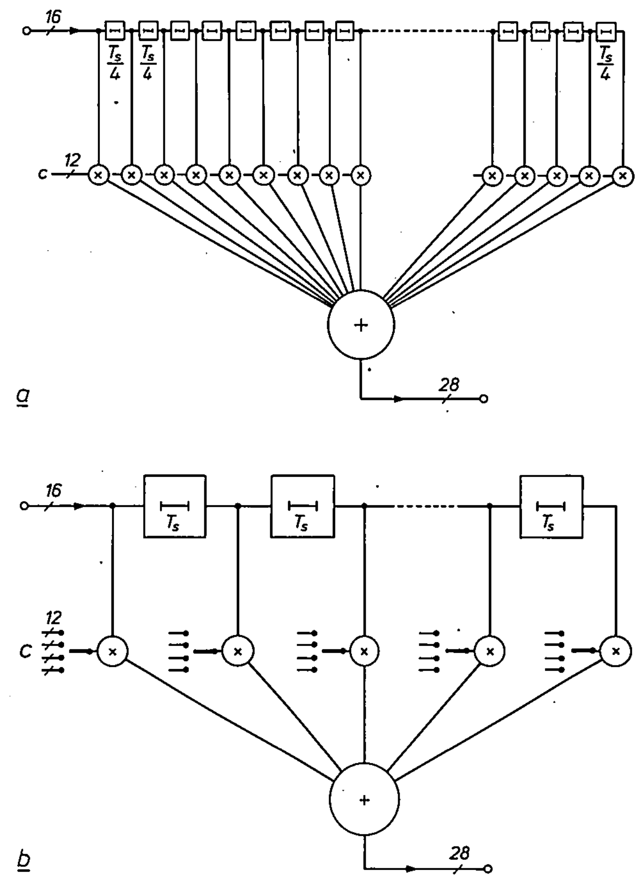

Fig. 4. Digital transversal filter. *a*) Filter consisting of 96 elements. A 16 bit word remains in each element for a quarter of the sampling period T_s. Since a new 16 bit word is only offered once per T_s, three-quarters of the elements are filled by the value zero. During the period T_s there are four multiplications by the 96 coefficients c; only 24 multiplications produce a product different from zero. These products are summed; in this way an output is provided four times in each sampling period, i.e. at a frequency of 4×44.1 kHz = 176.4 kHz. This means that there is a fourfold oversampling. *b*) An equivalent circuit that has been used in practice instead of (*a*) because it has 24 delay lines and multipliers instead of 96.

The practical version of the filter is in fact somewhat different from the version referred to in the above explanation. In practice the filter consists of only 24 delay elements and a 16 bit word remains in each element for a time T_s (fig. 4*b*). During this time T_s the word is multiplied four times by a coefficient c, which is different for each multiplication. The products are also summed four times during the time T_s and passed to the output. The frequency at which these summated values appear at the output is therefore $4/T_s$ = 176.4 kHz) again.

The coefficients are numbers with 12 bits. Each product has a length of 16 + 12 = 28 bits. The numbers have been chosen in such a way that the summation of 24 products does not introduce extra bits, so that the filter output consists of 28 bits with no rounding off.

The frequency spectrum of the oversampled and filtered signal is shown in fig. 3*c*. It can be seen that the bands in this spectrum around 1×, 2× and 3×44.1 kHz are suppressed.

The digital-to-analog converter generates a current whose magnitude is proportional to the last digital linear phase characteristic in the passband. This filter is simple and requires no highly accurate elements.

The hold function and the Bessel filter introduce some slight attenuation at the top of the passband. The digital filter is designed to correct this with a small overshoot (*fig. 5*).

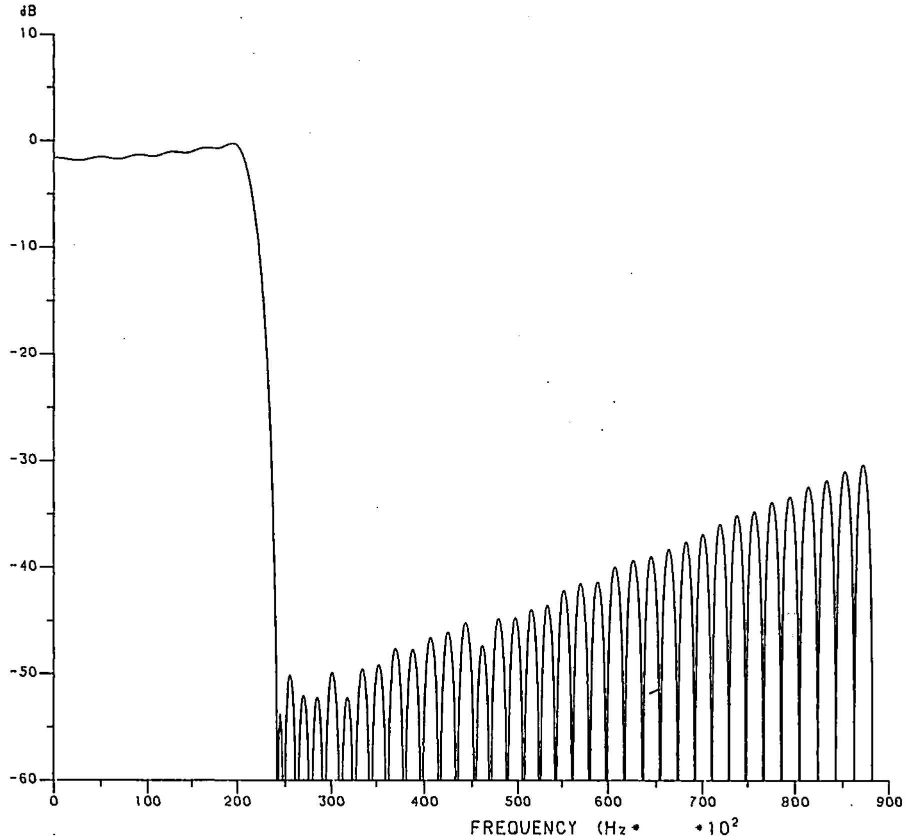

Fig. 5. Computer calculation of the detailed passband characteristic of the digital transversal filter. This has a small overshoot at the highest audio frequencies, which is used to compensate for the slight attenuation produced here by the curve in fig. 3*d* and the analog Bessel filter. A very sharp lowpass cut-off of 50 dB is obtained. The irregularity in the suppressed band is caused by rounding-off the filter coefficients to 12 bits.

word presented. This current is kept constant in a hold circuit until the next sample value is delivered, producing the staircase curve mentioned above. The signal samples have thus in theory changed from infinitely short pulses to pulses with the duration of a sampling period. This also has consequences for the frequency spectrum; the spectrum in fig. 3*c* is multiplied by a curve of the form $|(\sin x)/x|$ that has a first zero at 176.4 kHz (see fig. 3*d*). This gives an attenuation of signals in the 20 kHz sidebands on either side of 176.4 kHz by more than 18 dB. The hold effect causes no phase distortion.

The attenuation is still not sufficient, however. As a supplement, a lowpass Bessel filter of the third order is used, which has its −3 dB point at 30 kHz. The Bessel type of filter has been selected because of its

Suppression of the quantization noise

The presented signal, quantized to 16 bits, will contain some noise on conversion into an analog signal. This reproduces the errors due to the quantization in fixed steps. The root-mean-square value of the noise voltage in the sampled frequency band is $q/\sqrt{12}$, where q represents the magnitude of the quantization step. We see that when the quantization step is doubled, i.e. coding with one bit less, the noise voltage is also doubled, or, in other words, the noise level rises by 6 dB.

The samples that leave the filter at a repetition frequency of 176.4 kHz describe a signal with a bandwidth of 88.2 kHz. The quantization noise added due to the subsequent rounding off to 14 bits is spread over this band. With a signal of sufficient amplitude

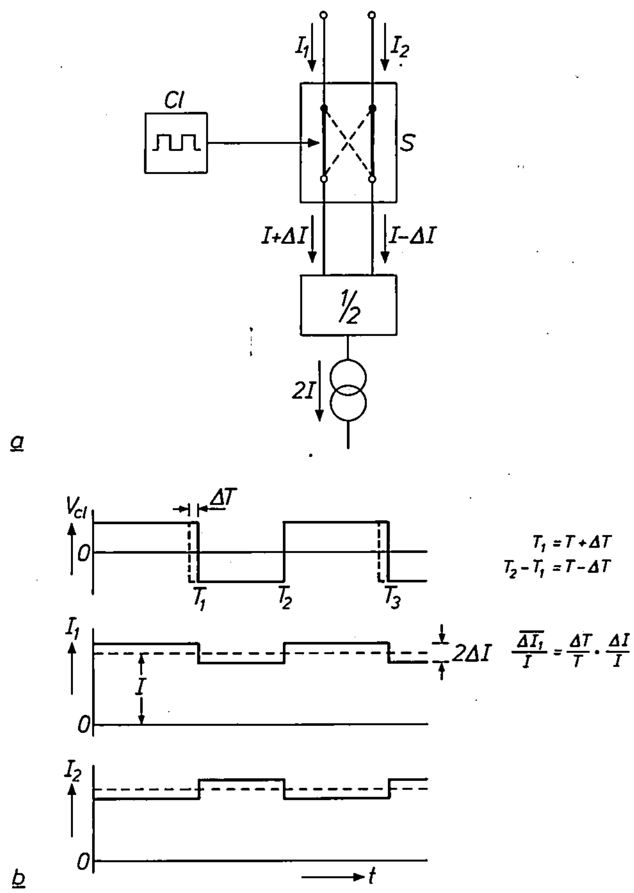

Fig. 6. *a*) Division of a current $2I$. *Cl* clock generator. *S* switches for periodically interchanging the two half-currents. *b*) The output currents I_1 and I_2 as a function of time *t*. Their mean value is the same. A difference between the mean output currents can be caused by an asymmetry ΔT of the clock signal V_{cl}. This difference is however an order of magnitude smaller than ΔI.

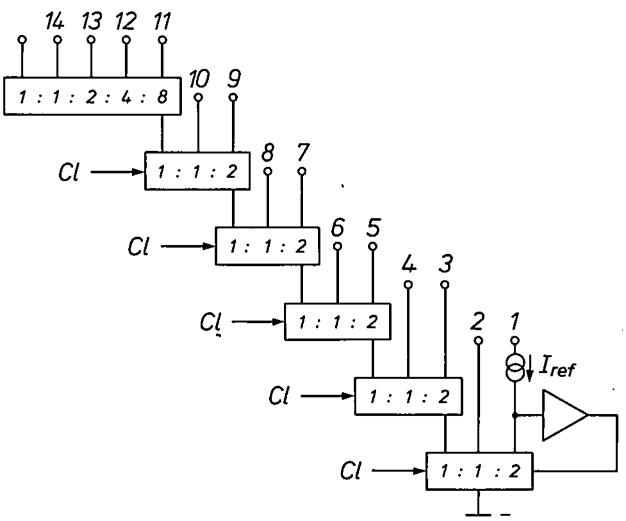

Fig. 7. Cascade of current dividers in the 14 bit digital-to-analog converter TDA 1540. The starting point is the reference current I_{ref}. Currents that are accurately equal to a half and a quarter of the input current are obtained in the divider stages by periodic interchanges; the clock signal *Cl* controls these interchanges. Only the four least-significant bits *11* . . . *14* are obtained by passive division.

and a sufficiently broad frequency spectrum this distribution is uniform, since the quantization errors for successive samples are in principle uncorrelated; the quantization noise is 'white' noise. Only the band from 0 to 20 kHz is relevant; this is only about a fourth part of the sampled band, and the noise power in the band from 0 to 20 kHz is therefore only a fourth part of the total noise power. This means that because of the fourfold oversampling the signal-to-noise ratio in the relevant frequency band is 6 dB better than would be expected with 14 bit quantization. It is thus about 90 dB, which is what would have been obtained with a 15 bit system without oversampling.

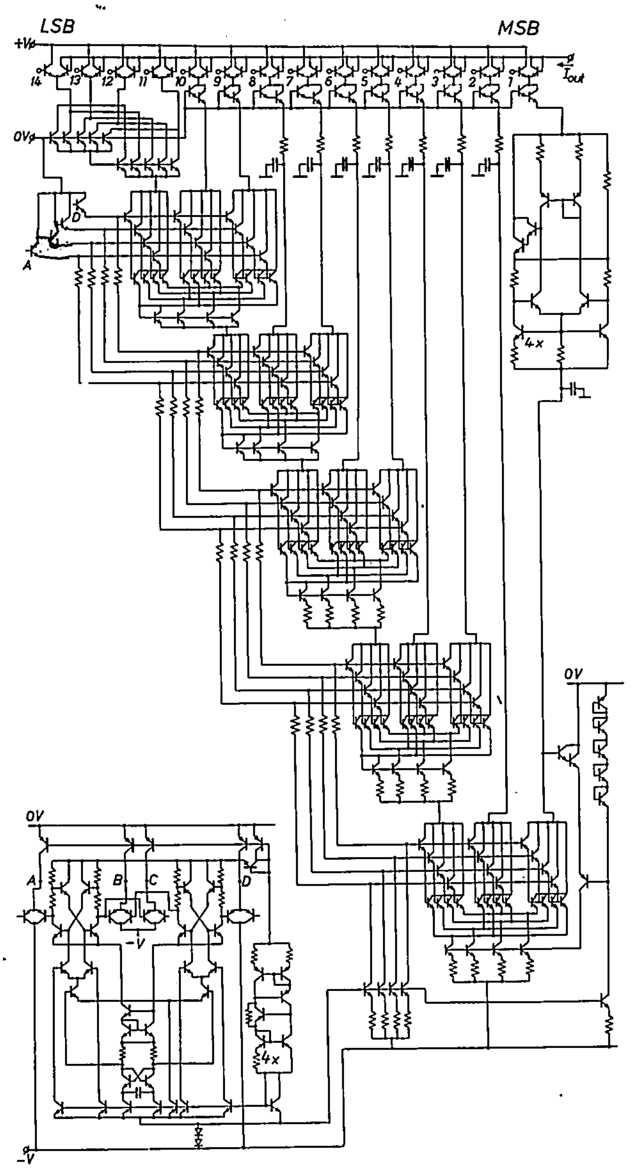

Fig. 8. Complete circuit diagram of the 14 bit digital-to-analog converter. The cascade of current dividers in fig. 7 can be identified here. The capacitors (*above*), which smooth out the ripple on the divider-output currents, are external. *Bottom left:* The clock generator.

In rounding off from 28 to 14 bits it is useful to compare successive rounding-off errors. If the analog signal is a direct voltage, successive samples will have the same rounding-off error. The audio signal will not contain any direct current; it will however contain slowly varying signals that will resemble a direct current in a short time interval. If the error produced in the rounding-off from 28 to 14 bits is now changed in sign and added to the next sample to arrive (see fig. 2), the average quantization error for slowly varying signals — i.e. low frequencies — can be reduced. This appears in the shape of the frequency spectrum of the quantization noise (see fig. 3*e*); at low frequencies the noise level is lower, at high frequencies it becomes higher. With a sampling rate of 176.4 kHz, it follows that a 7 dB gain in signal-to-noise ratio is obtained in the audio band (0-20 kHz). The ratio of the maximum signal to the noise contributed by the entire digital-to-analog conversion system described above is thus brought to about 97 dB, i.e. the value corresponding to a 16 bit quantization.

The digital-to-analog converter

The 14 bit digital-to-analog converter has been dealt with in detail elsewhere [2]. Here we shall only indicate how it differs from other digital-to-analog converters.

A characteristic feature is the way in which currents are generated that are accurately related by a factor of 2; a digital-to-analog converter requires a set of such currents. The exact ratio is obtained by periodically interchanging the currents that are derived by dividing down by two from a constant reference current (see *fig. 6*), so that small differences are averaged out. This system is known as 'dynamic element matching'. Accurate division by four can be carried out with a slightly more complicated circuit, also based on periodic interchange. The full series of current dividers is shown in *fig. 7*. Here *Cl* is the clock signal that controls the periodic switching; only for the four least-significant bits are the currents obtained from a passive division by means of differences in emitter area.

Fig. 8 shows the complete switching diagram of the 14 bit digital-to-analog converter. The cascade of divider stages can be seen in the figure. The ripple caused by the periodic switching is smoothed at the seven most significant bits by an *RC* filter; the seven capacitors (above in fig. 8) are externally connected.

The nonlinearity of the digital-to-analog converter is extremely low: between $-20\,°C$ and $+70\,°C$ it is less than 3×10^{-5}, or half the least-significant bit. The TDA 1540 integrated circuit is followed by the lowpass Bessel filter of the third order, and the analog signal appears at the output.

[2] R. J. van de Plassche and D. Goedhart, A monolithic 14-bit D/A converter, IEEE J. SC-14, 552-556, 1979.

Summary. The 16 bit words from the error-correcting circuit are converted into an analog signal by a 16 bit conversion system. This system consists of a digital transversal filter, in which the signal is oversampled 4 times (sampling rate 176.4 kHz) and then filtered in such a way that signals at frequencies above 20 kHz are attenuated by 50 dB after digital-to-analog conversion. The filter is followed by a noise shaper, which rounds off to 14 bits with negative feedback of the rounding-off error of the preceding sample. Next there is a 14 bit digital-to-analog converter, which is followed by a lowpass third-order Bessel filter. The signal-to-noise ratio of the complete system is about 97 dB. Even though the lowpass filter has a sharp cut-off the system is phase linear. The entire system, except for a few operational amplifiers, is contained in three integrated circuits; one for the digital filter (for both of the stereo channels) and two for the two digital-to-analog converters.

The Compact Disc has a diameter of only 12 cm. The information is recorded on the side of the disc shown here, protected by a transparent layer; the other side carries the label. The high density of the information gives a continuous playing time of more than an hour. The disc is packed in a 'de luxe' case, which also contains programme information.

Volume 40, 1982, No. 6 pages 149-180 Published 10th September 1982