CERN/ECP 91–5
26 March 1991

POSSIBLE APPLICATIONS OF THE SIGMA DELTA DIGITIZER IN PARTICLE PHYSICS

B. Hallgren

CERN, Geneva Switzerland

ABSTRACT

The sigma delta ($\Sigma\Delta$) principle is an analog-to-digital conversion technique based on high-frequency sampling and low-pass filtering of the quantization noise. Resolution in time is exchanged for that in amplitude so as to avoid the difficulty of implementing complex precision analog circuits, in favour of digital circuits. The approach is attractive because it will make it possible to integrate complete channels of high resolution analog-to-digital converters and time digitizers in sub-micron digital VLSI technologies. Advantage is taken of the fact that the state-of-the-art VLSI is better suited for providing fast digital circuits than providing precise analog circuits. The paper describes the principle and the performance of the ideal $\Sigma\Delta$ digitizer. The design and measurements of a new 10 MHz prototype circuit of a second-order $\Sigma\Delta$ is presented to show the high speed operation of such a circuit. The expected performance of a CMOS test design using the same principles is discussed. Digital filters, useful for particle physics, are introduced. A comparison to other digitizing techniques are made and the potential applications of the $\Sigma\Delta$ digitizer in particle physics are outlined.

Submitted to Nuclear Instruments and Methods in Physics Research, section A.

Mac/5021F/BH/ef

1. INTRODUCTION

The ΣΔ digitizing technique has been used in the telecommunications industry for many years. Its applications have been restricted to fairly low frequencies, such as the audio spectrum. Nevertheless, there is no fundamental reason why the ΣΔ digitizer could not be used at much higher frequencies by employing current VLSI technologies. The name ΣΔ is explained by the basic operation performed in the converter: a recursive integration of the difference between the signal current and a digital feedback current. Alternatively the name with transposed characters ΣΔ is sometimes used to identify the same equivalent circuit(*) [1]. There exist several waveform digitizing systems under the generic name delta modulation (ΔM), not to be confused with ΣΔ. In comparison to these methods the ΣΔ is particularly well suited for particle physics implementations because it integrates the input signal and thus can measure signal charge directly.

The single integration ΣΔ digitizer is one of the simplest ADCs known, since it can be realized as a one-bit waveform digitizer. Hence, it requires only one set of the basic ADC components: the comparator and the D flip–flop. The ΣΔ uses oversampling (sampling the input signal at hundreds of times higher frequency than is motivated by the signal bandwidth) to compensate for the lowresolution ADC. Digital signal processing is needed to extract the information and convert the 1-bit output data stream to a multi-bit digital word, similar to that of a flash-ADC output. This can be done with an algorithm implemented either in a real-time (fig. 1(a)) high-speed digital hardware circuit which may or may not be programmable, or with a subroutine in an off-line computer (fig. 1(b)). The term digital filter is applied to the hardware or the specific software routine in execution. In particle physics, the second solution would appear appropriate since only the data of interest need to be filtered. Since the ΣΔ does not require precision components in its analog circuitry, it is possible to design an ADC/TDC pipeline onto one integrated circuit using a digital VLSI process. The chip would consist of multiple ΣΔ digitizers, each with its own digital memory of a few kilobits, and would result in a cost-effective front-end electronics solution. Such circuits could be mounted onto a detector, and read out digitally in a time-multiplexed mode.

The sacrifice for the simplicity in hardware of the single integration ΣΔ is not so much the need for digital signal processing, but rather the somewhat lower signal bandwidth. The concept of ΣΔ has therefore been extended to higher-order integration converters. The order of the ΣΔ is defined as the number of integrators and feedback loops used in the circuit. A zero order ΣΔ has no feedback and is simply an integrator followed by a one-bit flash ADC. Second-order and higher-order ΣΔ digitizers are more

---

(*) The z-transform methods are described in many digital signal processing text-books.

appropriate for situations which require resolutions of 10 or more bits. It is believed that it should be technically possible to accommodate signal shaping times suitable for the next generation of hadron colliders (LHC and SSC) with state-of-the-art CMOS technology.

2. THE ΣΔ PRINCIPLE

The first-order ΣΔ is shown in fig. 2. The input signal current is integrated on one capacitor C. A negative feedback control circuit consisting of a single-bit ADC and a current switch (DAC) regulates the voltage on the capacitor. The DAC has a current output with two values $+i_D$ and $-i_D$ and is connected to the inverting output of the D–F/F. The operation of the circuit is the following: if at the clock transition the voltage on C is larger than zero the ADC output is logic one, while if the voltage is below zero the digital output is zero. A logic zero input to the DAC should result in $-i_D$ as feedback current and a reduction of the capacitor voltage. As a first example, (fig. 2(b)) assumes that there is no input signal and that the voltage on C is slightly positive before the clock edge. The ADC then makes the decision so that the feedback current is $-i_D$ during the clock period and therefore C constantly discharges in negative direction to a voltage called $-v_n$. At the next clock edge the feedback current will change polarity. The voltage on C will rise and reach the same slightly positive value as at the first clock transition, and the procedure will repeat again. The digital output then consists of alternating zeros and ones (10101010).

In the second example (fig. 2(b)), there is a constant input signal with the value equal + 1/3 i_D applied at just before the clock with the same initial conditions as above. The voltage on C will for this case reach – 2/3 v_n and then at next clock period + 2/3 v_n and thereafter the slightly positive initial value. The digital output from the ADC will have an extra one (101101101101) in the bit pattern. The interpretation of the bit stream is done by altering the digital word representation 1 and 0 to the values ± 1 and taking the average value over n samples (n is typically in the range 50 to 500 for a first-order ΣΔ). Although more efficient, digital filtering methods will be described later. Ideally the filtering time (n times the sampling clock period) should be matched to the charge collection time of the physics detector. In that case there will be no dead time or other speed disadvantage using the ΣΔ.

A detailed analysis of the above circuit will show that the accuracy of the conversion will not be constant but strongly dependent on the amplitude of the input signal. As an example of this, the value 1/3 i_D is particularly easy for the circuit to find, while any value which is not an integer fraction of feedback current will be longer to convert into a bit stream value. In a practical circuit a high-frequency dither signal is added on the integrator to ensure that there is equal probability for an input signal of any value within the full-scale range ± i_D to converge while satisfying the requirements of precision and

conversion time. The dither signal is eliminated from the digital output by low-pass filtering.

2.1 The first-order ΣΔ converter

The common method of analyzing theΣΔ circuit is with help of the z transform method. The z transform plays a role in discrete-time system theory analogous to that of the Laplace transform in continuous- time system theory. The following explanation of ΣΔ is based on Candy [2]. A simplified discrete-time model of the ΣΔ is shown in fig. 3. The z^{-1} is the delay of one period. The integrator is implemented as a recursive sum, while the ADC/DAC quantization noise or error has been modelled by an additive white noise source called **e**. The noise **e** is assumed to be uncorrelated with the input signal. This is achieved with help of the dither signal. The r.m.s. value of **e** is given by eq. (1) valid for a uniform quantizer with the quantizing step i_D.

$$e_{\text{r.m.s.}} = \frac{i_D}{\sqrt{12}} . \tag{1}$$

The r.m.s. spectral density e(f) of the quantization noise is equally distributed in frequency from DC to half the sampling frequency $f_s/2$ due to the assumption that it is white. As the square-root of the integral over $f_s/2$ to 0 of $e(f)^2$ should be equal to $e_{\text{r.m.s.}}$, it follows that e(f) has the magnitude

$$|e(f)| = e_{\text{r.m.s.}} \sqrt{\frac{2}{f_s}} . \tag{2}$$

The understanding how the ΣΔ modifies e(f) as a function of frequency, is of interest. In fig. 3 it is shown that the ADC-noise **e** is multiplied by the factor $(1-z^{-1})$, which is the z-transform for a differentiator. The frequency response is found by replacing z^{-1} with $e^{-j\omega/fs}$; ω is the angular frequency 2 πf and f_s is the sampling frequency applied to the ADC. The spectral density of the noise η(f) from a first-order ΣΔ is therefore

$$\eta(f) = e(f)\,(1-z^{-1}) = e(f)\,(1-e^{\frac{-j\omega}{f_s}} . \tag{3}$$

From eqs (2) and (3), and with some algebra, the r.m.s. value of the spectral density of the first-order ΣΔ quantization noise is found to be

$$|\eta(f)| = e_{\text{r.m.s.}} \sqrt{\frac{2}{f_s}}\; 2 \sin \frac{\omega}{2 f_s} . \tag{4}$$

The eq. (4) illustrates the ΣΔ principle; **it shapes the normal quantization noise spectrum of the ADC.** The negative feedback makes the resolution of the ΣΔ frequency dependent. This is explained by the fact that the coarse digital feedback signal

is averaged by the analogue integrator before the ADC. The longer time the integrator can treat the feedback signal, the better the approximation of the real signal will be. Thus, low-frequency signals are better measured than high frequency.

In fig. 4 both the flat noise spectrum of an unmodified ADC with a 1-bit resolution (eqs (2) and (4)) is shown. The two spectra cross at the frequency equal to $f_s/6$. Above this frequency the effect of the digital feedback is to make the noise worse. A zero-order ΣΔ will be the most effective choice for an improvement of 1 to 1.5 bits. In this case only the digital filtering of the bit stream has to be implemented to do signal averaging.

The resolution of the first-order ΣΔ in signal-to-noise ratio (SNR) and effective bits is given for the three bandwidths in fig. 4. A sinewave signal with the amplitude equal $i_D = 1.0$ (full scale) is used for calculating the SNR. The r.m.s. value is $1/\sqrt{2}$ and the frequency of the signal is $f_s/1000$. The r.m.s. value of the noise is obtained by integrating the noise spectrum from DC to the frequencies $f_s/128$, $f_s/64$ and $f_s/32$, respectively. Note that 1.5 bits in resolution is gained by every decrease in bandwidth by a factor 2. This is 1 bit more than by signal averaging only. An ideal low-pass filter with brick-wall characteristics was used in fig. 4. The eq. (5) [3](*) shows the theoretical improvement over the original 1-bit in the resolution r for a first-order ΣΔ for realistic filters with the bandwidth f_b.

$$r = 1.5\ n - 1.2,\ n = \log_2 \left(\frac{f_s}{2 f_b}\right). \tag{5}$$

The eq. (5) indicates that a first-order 1-bit ΣΔ clocked at the frequency 256 MHz followed by a filter with 2 MHz bandwidth can measure signals below this frequency with 7.8 + 1 bits accuracy.

2.2 The second-order ΣΔ converter

By making the negative feedback more effective it is possible to increase the resolution of a ΣΔ system. The second-order ($\Sigma^2\Delta$) consists of two feedback loops and two integrators, which are connected by a transconductance amplifier GM with current output as shown in fig. 5. The digital output are still one-bit words. This circuit can be analyzed using the same techniques as for the first-order ΣΔ. The result of eq. (6) shows that the quantization noise is multiplied by the same factor as in the first-order ΣΔ, but squared.

(*) The proceedings from the IEEE International Symposium on Circuits and Systems (ISCAS) of recent years contain many papers on the ΣΔ subject. The ISCAS'90, May 1–3, 1990 had 18 ΣΔ papers on this subject.

$$|\eta(f)| = e_{\text{r.m.s.}} \sqrt{\frac{2}{f_s}} \left(2 \sin \frac{\omega}{2 f_s}\right)^2 . \qquad (6)$$

Figure 6 shows the much increased efficiency of the $\Sigma^2\Delta$ to suppress the quantization noise in comparison with the first-order $\Sigma\Delta$; a resolution of 2.5 bits is gained by every factor of two decrease in bandwidth as shown in eq. (7)

$$r = 2.5\ n - 2.5,\ n = \log_2 \left(\frac{f_s}{2 f_b}\right) . \qquad (7)$$

The eq. (7) shows, for example, that a 1-bit $\Sigma^2\Delta$ clocked at the frequency 256 MHz followed by a filter with 8 MHz bandwidth will have a resolution of 8.5 bits.

2.3 Simulations and measurements

Functional simulating techniques are used to do the detailed design of a $\Sigma\Delta$ system because conventional circuit and systems analysis methods (such as SPICE) have, to date, proved to be intractable. Also evaluations of prototype circuits are useful in understanding the $\Sigma\Delta$ principle. Figure 7 shows the circuit diagram of a second-order $\Sigma\Delta$ converter. It consists of two integrators made of differential transistor pairs and 510 pF, respectively 330 pF capacitors. The ADC is made with one comparator and a D–F/F, while the DAC is a current switch transistor pair. The digital output is connected to an available CAMAC memory module LeCroy 4302. Figure 8 shows a typical measurement on the 10 MHz second-order $\Sigma\Delta$ of fig. 7 and table 1 outlines the results from the measurements on this second-order $\Sigma\Delta$. The digital output from the circuit was directly used as input to a FFT program. The main difference between the experimental circuit and the theory using ideal components is the noise floor. However, if in the functional simulator the finite low-frequency gain of the integrators in fig. 7 and the propagation delay (10 ns) of the ADC is taken into account, the measurements agree with the simulations within 0.5 bits.

The useful signal range, or dynamic range (DR), of the $\Sigma\Delta$ for sinusoidal inputs is defined as the ratio of the full-scale input signal to the input signal at which the SNR is unity (0 dB). Figure 9 shows the DR results of measurements on the same second-order $\Sigma\Delta$ but with different gains in the differential amplifier.

The gain can be changed by the resistor (RE), (fig. 7). The curves show the performance for RE = 10 Ω, 100 Ω and 1 kΩ, respectively. The curve for 1 kΩ shows a DR of 1.7 V to 0.5 mV = 3400/1 or 11.4 bit. The curves for, especially, 10 Ω shows the well known linearity problem of a differential amplifier with low emitter resistance. However, for a 3 mV r.m.s. input signal it is possible to have more than 8 bit resolution. A signal splitting technique could be used to achieve a dynamic range of nearly

17 bits. The measurements were made with a digital filter with a cut-off frequency equal to $f_s/128$.

2.4 Expected performance of a CMOS test design

Several breadboard circuits similar to fig. 7 have been built with discrete components of both first-order and second-order ΣΔ which confirm the working of the ΣΔ at frequencies between 100 and 200 MHz. Based on these prototypes and a system level simulator, a prototype chip has been designed in a low-cost 2 μm CMOS n-well, double poly, double metal process. The main difference between the CMOS prototype and the theory, as described in sect. 2, is that the delay of the comparator and DAC is not negligible compared to the clock period. The delay is also dependent on the amplitude of the input voltage to the comparator.

Simulations based on typical process parameters show that the delay to be expected is ~ 2.2 ns for an input voltage of 5 mV. Compared to discrete components this is a good performance and is due to the fact that no output stage is needed to drive large currents. The delay of an inverter in the technology used is 0.4 ns. The simulations show that a maximum clock frequency of f_s = 160 MHz could be reached to be confirmed by measurements. The simulated performance agrees well with table 1, e.g. with a filter with a cut-off frequency equal to $f_s/128$ a dynamic range of 4000 to 1 could be expected. In a particle physics application using the filters described in sect. 3.3, the peaking time T_m should be multiplied with 6 ns. Therefore, table 2 suggests a 12-bit dynamic range with a peaking time of 700 to 900 ns and table 3 indicates an 8-bit performance with a peaking time of < 230 ns.

2.5 Other ΣΔ configurations

It is possible to increase the number of feedback loops in ΣΔ and thereby increase the slope of the quantization noise reduction with frequency. Both third-order and fourth-order ΣΔ have been built. These higher-order ΣΔs tend to be marginally stable but many solutions have been proposed to solve these problems [3]. The resolution r increases in bits as a function of the oversampling ratio for a third-order ($\Sigma^3\Delta$) is shown in eq. (8).

$$r = 3.5\ n - 3.7,\ n = \log_2 \left(\frac{f_s}{2 f_b}\right) . \qquad (8)$$

Equation (8) shows that a 256 MHz 1-bit ($\Sigma^3\Delta$) with a bandwidth of 8 MHz has 11.3 bits resolution.

3. FILTERS

3.1 Introduction

The digital data coming from the ΣΔ has to be converted into a useful format which is done by a digital filter. The filter should eliminate the hig-frequency quantization noise generated by the ΣΔ but also fulfil the normal tasks of a filter in a particle physics front-end: optimize the signal-to-noise ratio and the shape of the signal. The implementation of these filters using digital signal processing techniques will be described below. But before treating this subject, it is necessary to discuss the general problems caused by the effect of sampling a continuous signal [4]. The sampling of a detector signal at the frequency f_S will cause the frequency spectrum of this signal centred around DC (the frequency = 0) to repeat around each multiple of the sampling frequency: $\pm f_S$, $\pm 2 f_S$, etc. It can be seen that the two signal spectra centred at the frequencies 0 and f_S will overlap at the frequency $f_S/2$. The $f_S/2$ is called the folding frequency. It is important to keep the total power of signals above $f_S/2$ smaller than the required minimum resolution so that they are not spuriously mixed. The unwanted mixing phenomenon is called aliasing. The aliasing may occur in particle physics waveform recorders such as flash-ADC systems and analogue-pipeline systems.

3.2 The anti-aliasing filter

There is little that can be done to remove aliased power once it has been discretely sampled. Therefore, an analog anti-aliasing filter is needed to suppress sufficiently all frequency components of signals higher than the folding frequency before they are sampled. These components consist of the high-frequency spectral components added together from various sources such as the detector signals and noise, from noise in the preamplifier and noise pickup from power supplies and external sources. The ambition when designing a robust digitizing system for LHC should be to make the attenuation of the filter at $f_S/2$ and above equal to at least the dynamic range of the digitizer system.

The frequency characteristics of the filter have to be such that it produces a clean time response to minimize pile-up. A filter with a very sharp frequency cut-off will have a time response with long damped oscillations. This is usually not acceptable in particle physics applications, hence the use of CR–RC filters in shapers. The following example illustrates this point: a single CR differentiation and three RC integrations with 33 ns time constant will give a signal which peaks at 100 ns. The corresponding frequency response of this filter is found to be: the –3 dB bandwidth is at 5.2 MHz and at 100 MHz; the high-frequency attenuation is –69 dB sufficient for 11 bits dynamic range. With this filter a 200 MHz sampling frequency would be needed to fulfil the anti-aliasing filter requirements. A tenth-order Bessel filter (a standard filter whose transient response has less than 1% overshoot) could reduce the necessary sampling frequency to ~ 25 MHz,

but such high-order filters are very expensive because they need up to tensets of high-accuracy components and amplifiers.

The great advantage of the ΣΔ systems is that they typically employ large oversampling ratios. Therefore, RC filters with one or two sections will often have sufficient attenuation at the foldover frequency.

3.3 The digital filter

A real time digital hardware filter can be made with only simple arithmetic operations: addition, subtraction and delays, instead of analog circuit elements like resistors, capacitors and amplifiers. The development of the VLSI technology entirely motivates this choice.

Most ΣΔ systems use one type of filter [6], which is the digital equivalent to "the clipped delay line filter" frequently used in physics. One section is shown in fig. 10. These filters are known under several names: the recursive running sum (RRS), the SINC filters or the COMB filter. Equation (9) shows the recursive formula for the low-pass filter of fig. 10(a).

$$y_n = y_{n-1} + \frac{1}{N}(x_{n-1} - x_{n-1-N}) . \qquad (9)$$

The input data is x_n and the output result is y_n. Figure 10(b) shows a high-pass filter consisting of a low-pass filter section with one unit delay plus a subtracter. The recursive formula of the high-pass filter is shown in eq. (10).

$$y_n = y_{n-1} + x_n - x_{n-1} - \frac{1}{N}(x_{n-1} - x_{n-1-N}) . \qquad (10)$$

The design of a digital shaper can be made with three low-pass sections and one high-pass. It is convenient to choose the delay length (N) in powers of two such as 16, 32 and 64, because the gain of the filter can then be kept to unity by division with these numbers which is implemented with shift operations. Figure 11 shows the time response of the filter for different lengths 16, 32 and 64 of the high-pass filter for an input step function. The shape of the response is similar to an analogue CR–RC chain [4] and very suitable for the applications described here.

Two examples illustrate the use of the digital shaper. The first one is a high-resolution 12-bit calorimeter type application; the second is an 8-bit flash equivalent application. In both cases it is assumed that the ΣΔ converter is connected to a charge preamplifier, the output of which is a step function. The total noise spectrum from a ΣΔ is simulated with the measured FFT spectrum of fig. 8, but with the sinewave signal and distortion components replaced by the average noise level in neighbouring frequenc bins.

This FFT spectrum has been obtained by weighting the sampled signal using a Hanning window function. For the high-resolution application the low-pass sections LP1, LP2 and LP3 are used with the lengths equal to 64. Three different lengths HP equal to 64, 128 and 256 have been chosen for the high-pass filter. In table 2 the values are given for the overall noise reduction (NR) defined as the ratio of the total noise without and with filter, the peaking time T_m (in sampling periods) and the amplitude reduction A_m (of full scale) or the gain of the filter. The noise reduction is about the same for the three cases. It can be seen that the high-pass filter has a large effect on the amplitude and therefore determines the dynamic range. The product DR = NR * A_m is given for a bipolar signal (sinewave) with the peak amplitude ± A_m. For low-rate particle physics applications with unipolar signals this figure could be doubled. The low-resolution filter example has three low-pass sections LP1, LP2 and LP3 each having 16 clock periods in delay, while the high-pass section (HP) has either 16, 32 or 64. In table 3 the results from calculations are shown.

Tables 2 and 3 can be used to predict the performance of a complete system using the ΣΔ. The example uses a ΣΔ front-end running at 1 GHz sampling clock. A 12-bit digitizer employing the high-resolution filter with sections 128, 64, 64, 64 will give pulse shapes peaking at T_m = 134 ns. A low-resolution 8-bit digitizer can be built for a peaking time of 34 ns with the 32, 16, 16, 16 filter. The response would resemble fig. 11 with one time step along the x-axis equal to 1 ns.

4. COMPARISON WITH THE TELECOMMUNICATIONS ΣΔ ADC, THE FLASH ADC AND ANALOG PIPELINES

4.1 The telecommunications ΣΔ ADC

It should be noted that the operation of the ΣΔ digitizer, as introduced in this paper, differs from the methods used in the telecommunication industry because in these the digital output is decimated by keeping only every nth word (with n equal to the oversampling ratio). In physics applications all digital words after the filtering may be utilized. This modification increases the usefulness of the ΣΔ digitizer because it can be a TDC and at the same time an ADC analogous to the flash ADC measuring time and amplitude.

4.2 The flash ADC

The ΣΔ is a circuit realized with the strict minimum of components while the flash ADC is one the most complicated electronics circuit built. There are at least 2 orders of magnitude more comparators in a 8 to 10-bit flash ADC than in a ΣΔ, which has one. The ΣΔ has a current integrating input ideally suitable for physics and requires a simple anti-aliasing filter. The flash ADC has to have a complex shaping and anti-aliasing filter to

adapt it for particle physics signals. In a digital pipeline the amount of memory for a 1 GHz ΣΔ equals that of a 10-bits 100 MHz system. As it is possible to integrate the ΣΔ in one multi-channel VLSI circuit, this solution will represent a considerable saving compared to a 10 to 12-bits flash ADC multi-chip system.

4.3 Analog storage techniques

The analog pipeline uses capacitors to store the signal temporarily until it can be digitized with a slow ADC. This classical technique has been implemented with discrete capacitors, charge-coupled integrated circuits (CCD) and, recently, switched capacitor circuits (SC). The SC pipeline has, typically, one hundred 1 pF capacitors. These are connected to the input with a transistor switch. However, the input of a SC circuit is not signal current integrating. It has to have the same complex shaping and anti-aliasing filter as the flash ADC. Another drawback is that one input signal is put on several capacitors because it is usually longer than the sampling period and not synchronized with the clock. The storage on the capacitors must be done by turning-on and at the same time turning-off the input transistor switches in a very precise manner. The tolerance of the capacitors is between 0.1 to 1% on the same chip; but is typically 30% from chip to chip. Great care needs to be taken so that, during the time of the temporary storage before readout, there is no influence on the capacitor. The stored charge can change because of crosstalk to the input and via substrate and power supply coupling. Experience shows that a sensitive system such as analog pipelines needs extensivecalibration efforts [7].

5. LHC APPLICATIONS

5.1 Calorimeter applications

Today, many existing calorimeter applications use the charge-integrating type of ADC. These applications could probably be handled with a 200 MHz $\Sigma^2\Delta$ digitizer with a digital filter and peaking times below 1 μs. An input amplifier with current output would be needed to adapt the input of the ΣΔ to a 50 Ω or 100 Ω system. A fast trigger output can be made either in the traditional analogue way or by using the digital bit stream to drive a DAC made with a single current switch and a programmable current source on which a weighting coefficient can be put.

The typical readout electronic chain of a liquid argon calorimeter, except the preamplifier, could be replaced by a ΣΔ digitizer channel where the analog shaper is made with a digital filter.

Future accelerators such as LHC will require approximately ten times faster response times. ΣΔ digitizers could be a very attractive solution if the sampling frequency were pushed to 1–2 GHz and/or if third-order or fourth-order ΣΔ were used. Today's state-of-the-art 1 μm CMOS technology is about three times faster than the 2 μm

process, for which sampling rates of 200 MHz are expected. In addition, 0.5 – 0.8 μm CMOS technologies are now or soon going into production. Recently an experimental CMOS circuit in such a technology was reported to be working at 4.2 GHz [7]. Thus, 1 GHz sampling rate for a CMOS ΣΔ digitizer does not seem impossible, from the standpoint of available technology. Bipolar and gallium arsenide E/D technologies promise even higher sampling rates, although probably at higher power dissipations (bipolar) and lower integration densities.

5.2 Drift chamber applications

For the next generation of high-luminosity hadron colliders, alternative tracking detectors such as the multidrift tube and the microstrip gas avalanche chambers are being studied. Clearly a ΣΔ digitizer would have to be clocked in the range 1 to 2 GHz in order to provide 100 to 200 ps r.m.s. time measurements needed in such detectors.

6. SUMMARY

The major advantages of the ΣΔ compared to other digitizing techniques are, they:

- are of very simple digitizer topology, in particular, the absence of analog pulse shapers/filters means higher accuracy and less variation between channels;
- can easily be pushed to high resolutions, 12 bits or more; has a current integrating input so that the total signal charge can be measured on a capacitor;
- require very little chip surface area, typically less than 0.6 mm^2 for the integrator, comparator, flip-flop and DAC in a 2 μm digital CMOS process;
- allow simple intermediate data storage since it requires only a few kilobits of static memory which could be integrated with the digitizer;
- are relatively easily integrated onto a detector since it allows simple serial digital readout (maybe time-multiplexed over high-speed links);
- are of off-line digital filtering (or through a special processor at readout) gives analysis flexibility since several filters could be applied to the same data, e.g. one optimized for timing analysis and another for best charge resolution;
- are potentially very cost-effective because multiple channels could be integrated onto a single IC, with a single common clock input.

The main disadvantages are, they:

- are difficult to design and requires frequency domain analysis techniques for the design and special purpose simulator programs;

- cannot match the performance of 6-bit or 8-bit flash ADC in general high speed applications;
- require oversampling means needing a very high-frequency clock, typically 200 MHz to 2 GHz depending on the application.

However, this is not expected to have a major influence on the system performance because any noise generated by the clock is synchronous and far outside the signal bandwidth.

Acknowledgements

The author is indebted to Dr V. Hungerbuehler of LeCroy SA for his important contributions to this CERN and LeCroy project during the past three years. It is a pleasure to acknowledge the support from Mr W. LeCroy of LeCroy Corporation.

REFERENCES

[1] P. Horowitz and W. Hill, The art of electronics, Cambridge Univ. press, Cambridge (1980).

[2] J.C. Candy, A use of double integration in sigma delta modulation' IEEE Trans. Commun., COMM–33 (March 1985) 249–257.

[3] G.C. Temes, A comparative overview of sigma delta modulator architectures, Lecture notes at the Lausanne intensive course in sigma delta modulators (June 1989).

[4] L.R. Rabiner and B. Gold, Theory and applications of digital signal processing, Prentice Hall, Englewood Cliff, NJ, USA (1975).

[5] P. Franzini, Noise and pile-up in liquid sampling calorimeters, proc. of the Workshop on Experiments, Detectors and Experimental areas for the Supercollider (1987) 636–656.

[6] J.C. Candy, Decimation for sigma delta modulation, IEEE Trans. Commun., COMM–34 (January 1986) 72–76.

[7] A. Honma et al., Performance of the front-end signal processing electronics for the drift chambers of the Stanford large detector, SLAC–PUB–5316 (October 1990).

[8] H. Cong et al., Multigigahertz CMOS dual-modulus prescaler IC, IEEE JSSC 23 (October 1988) 1189–1194.

Table 1

Resolution (in bits)	Bandwidth
6.6	$f_s/32$
9.2	$f_s/64$
11.4	$f_s/128$

Table 2

HP	LP1	LP2	LP3	NR	T_m	A_m	DR
64	64	64	64	9781	118	0.352	3443
128	64	64	64	9007	134	0.569	5125
256	64	64	64	9306	149	0.742	6905

Table 3

HP	LP1	LP2	LP3	NR	T_m	A_m	DR
16	16	16	16	253	30	0.368	93
32	16	16	16	266	34	0.580	154
64	16	16	16	283	38	0.748	211

TABLE CAPTIONS

Table 1 Measured performance of a second-order ΣΔ.

Table 2 Digital shaper, high-resolution example.

Table 3 Digital shaper, low-resolution example.

FIGURE CAPTIONS

Fig. 1 Sigma delta digitizer with: (a) real-time filter and (b) off-line filter.

Fig. 2 (a) First-order ΣΔ.
(b) Timing diagrams of the first-order ΣΔ for two input signal currents.

Fig. 3 Discrete-time equivalent model of the first-order ΣΔ.

Fig. 4 First-order ΣΔ converter theoretical quantizition noise spectrum (2^{15} points).

Fig. 5 (a) Second-order ΣΔ.
(b) Discrete time equivalent model.

Fig. 6 Second-order ΣΔ converter theoretical r.m.s. quantizition noise spectrum (2^{15} points).

Fig. 7 Circuit diagram of a 10 MHz second-order ΣΔ digitizer.

Fig. 8 Measured noise spectrum of the 10 MHz second-order ΣΔ converter (FFT with 2^{15} points).

Fig. 9 The dynamic range of the 10 MHz ΣΔ converter prototype for different gain settings.

Fig. 11 Fast digital shaper filters with high-pass filters of length 16, 32 and 64, respectively (table 3). The horizontal axis equals 10 sampling periods/division.

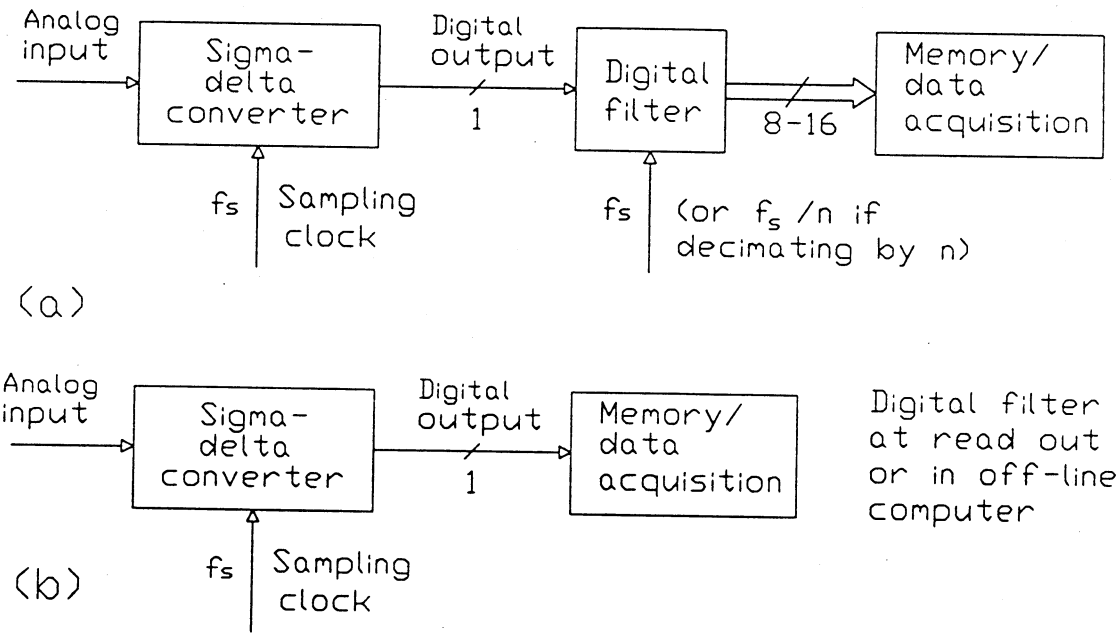
Analog input
Sigma-delta converter
Digital output
1
Digital filter
8-16
Memory/ data acquisition
fs
Sampling clock
fs
(or f_s /n if decimating by n)
(a)
Analog input
Sigma-delta converter
Digital output
1
Memory/ data acquisition
Digital filter at read out or in off-line computer
fs
Sampling clock
(b)

Fig. 1

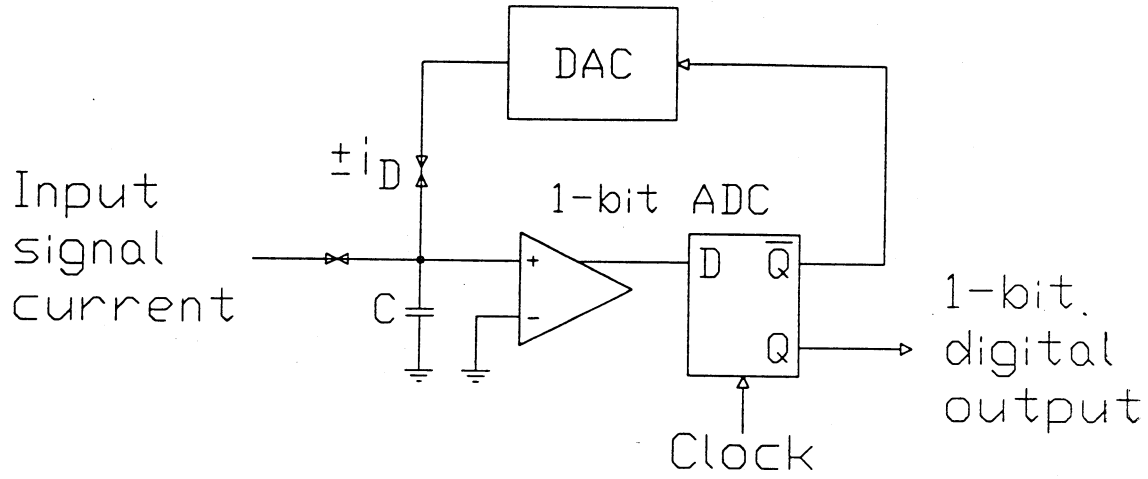
DAC
±iD
Input
signal
current
1-bit ADC
C
D
Q
Q
1-bit
digital
output
Clock

Fig. 2(a)

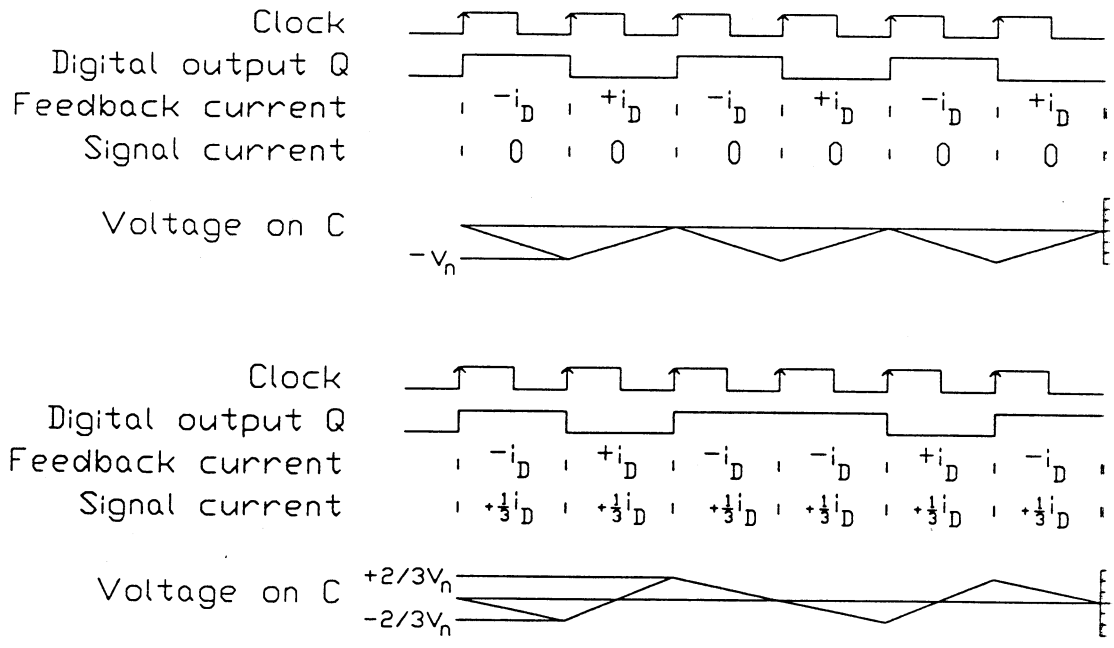
Clock
Digital output Q
Feedback current
Signal current
−iD
+iD
0
Voltage on C
−Vn
+2/3Vn
−2/3Vn

Fig. 2(b)

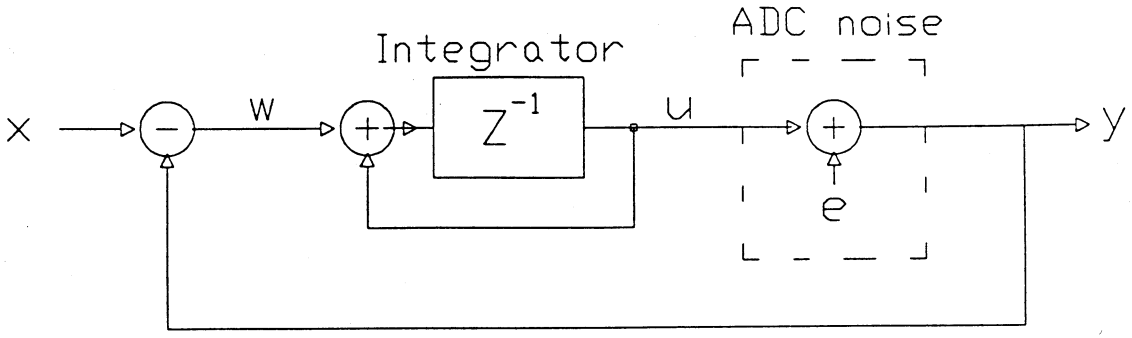

$$w = x - y$$
$$u = (w+u)\,z^{-1} \longrightarrow y = e(1-z^{-1}) + xz^{-1}$$
$$y = e+u$$

Fig. 3

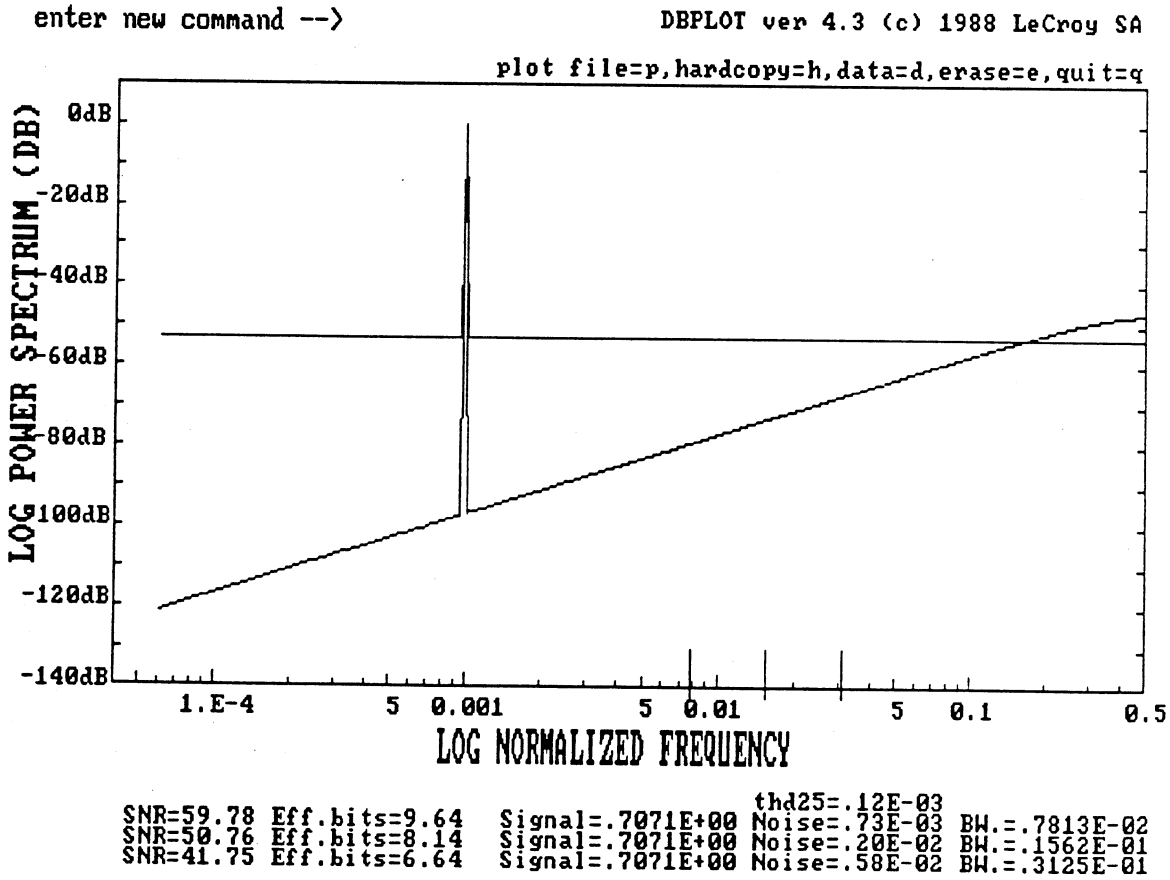

Fig. 4

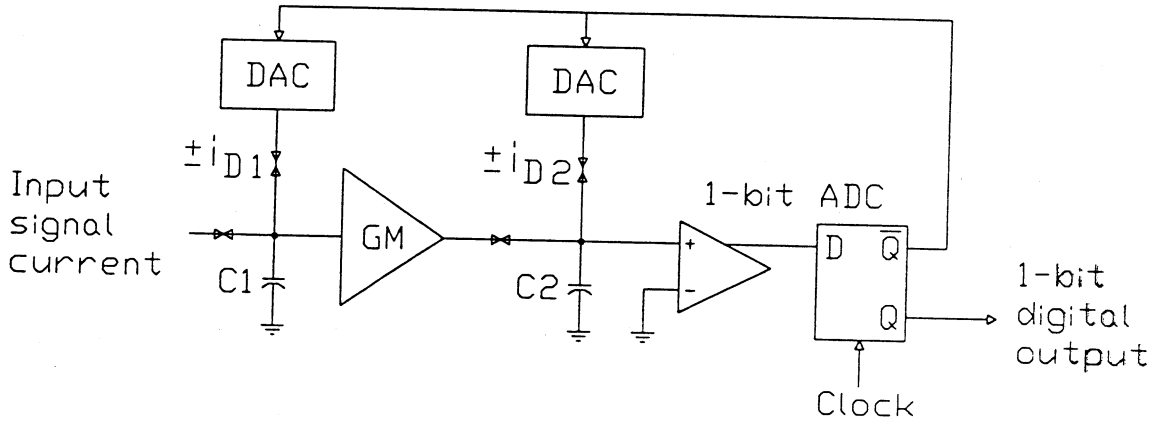

Fig. 5(a)

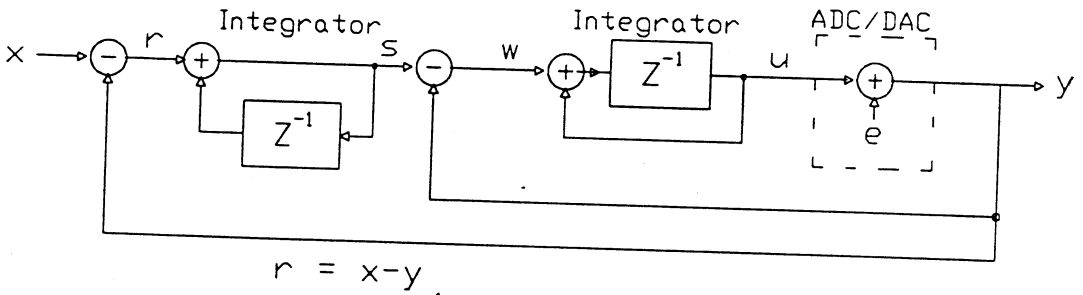

$$r = x - y$$
$$s = r + z^{-1}s$$
$$w = s - y \longrightarrow y = e(1-z^{-1})^2 + xz^{-1}$$
$$u = (w+u)z^{-1}$$
$$y = e+u$$

Fig. 5(b)

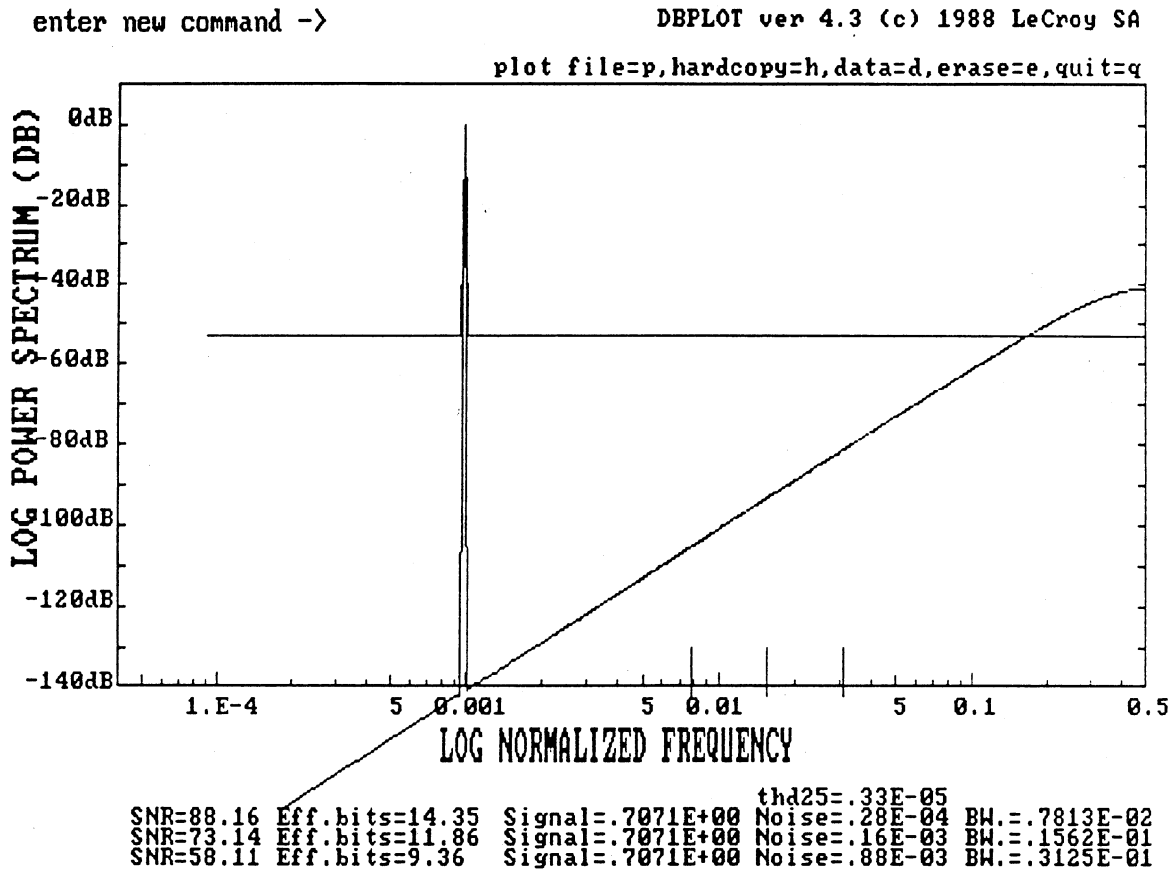
enter new command ->
DBPLOT ver 4.3 (c) 1988 LeCroy SA
plot file=p,hardcopy=h,data=d,erase=e,quit=q
LOG POWER SPECTRUM (DB)
0dB
-20dB
-40dB
-60dB
-80dB
-100dB
-120dB
-140dB
1.E-4
5
0.001
5
0.01
5
0.1
0.5
LOG NORMALIZED FREQUENCY
thd25=.33E-05
SNR=88.16 Eff.bits=14.35 Signal=.7071E+00 Noise=.28E-04 BW.=.7813E-02
SNR=73.14 Eff.bits=11.86 Signal=.7071E+00 Noise=.16E-03 BW.=.1562E-01
SNR=58.11 Eff.bits=9.36 Signal=.7071E+00 Noise=.88E-03 BW.=.3125E-01

Fig. 6

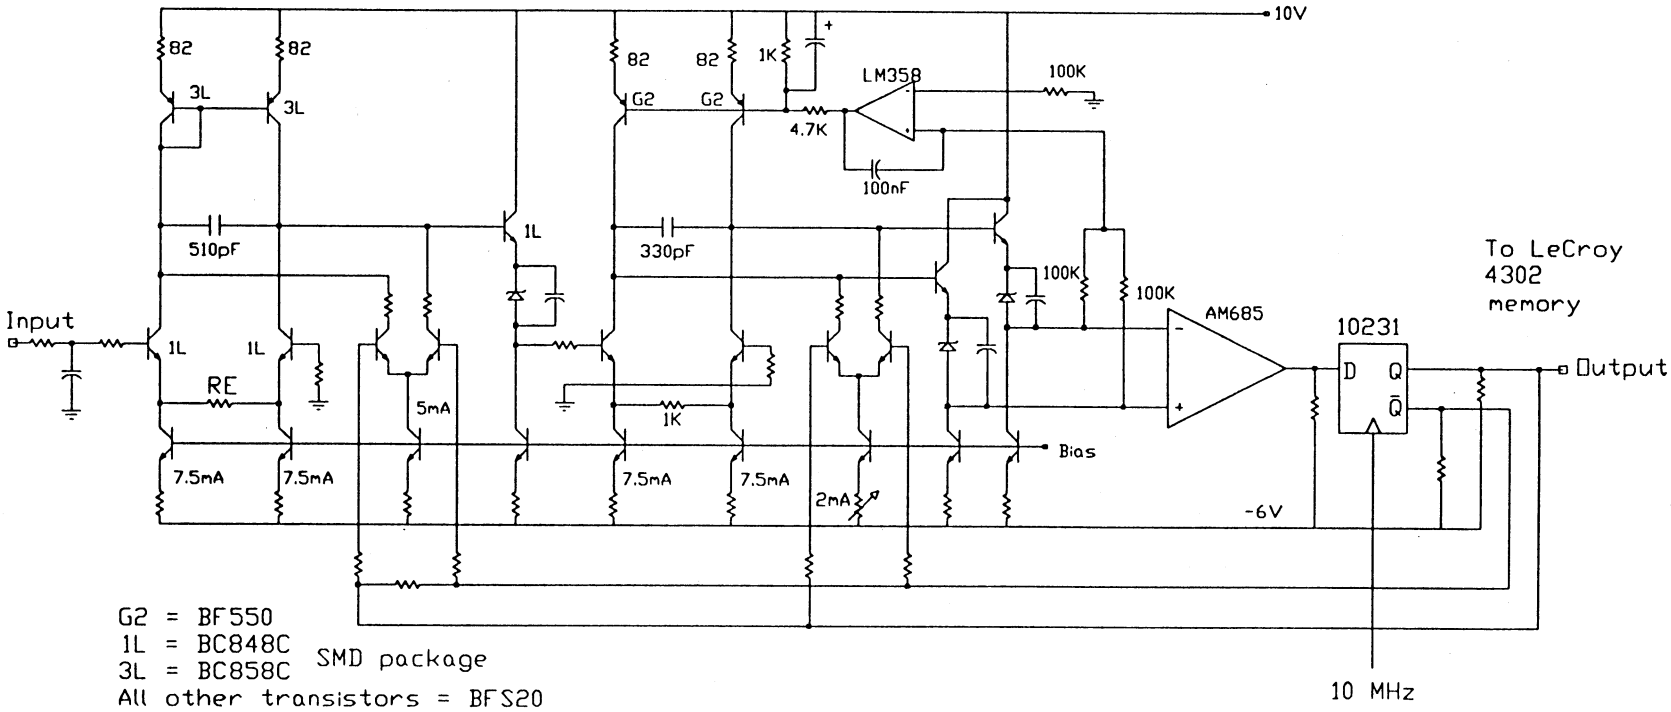
10V
82
82
3L
3L
82
82
G2
G2
1K
LM358
100K
4.7K
100nF
510pF
330pF
1L
100K
100K
AM685
10231
D
Q
Q̄
To LeCroy
4302
memory
Input
1L
1L
RE
5mA
1K
Bias
Output
7.5mA
7.5mA
7.5mA
7.5mA
2mA
-6V
G2 = BF550
1L = BC848C
3L = BC858C
SMD package
All other transistors = BFS20
10 MHz

Fig. 7

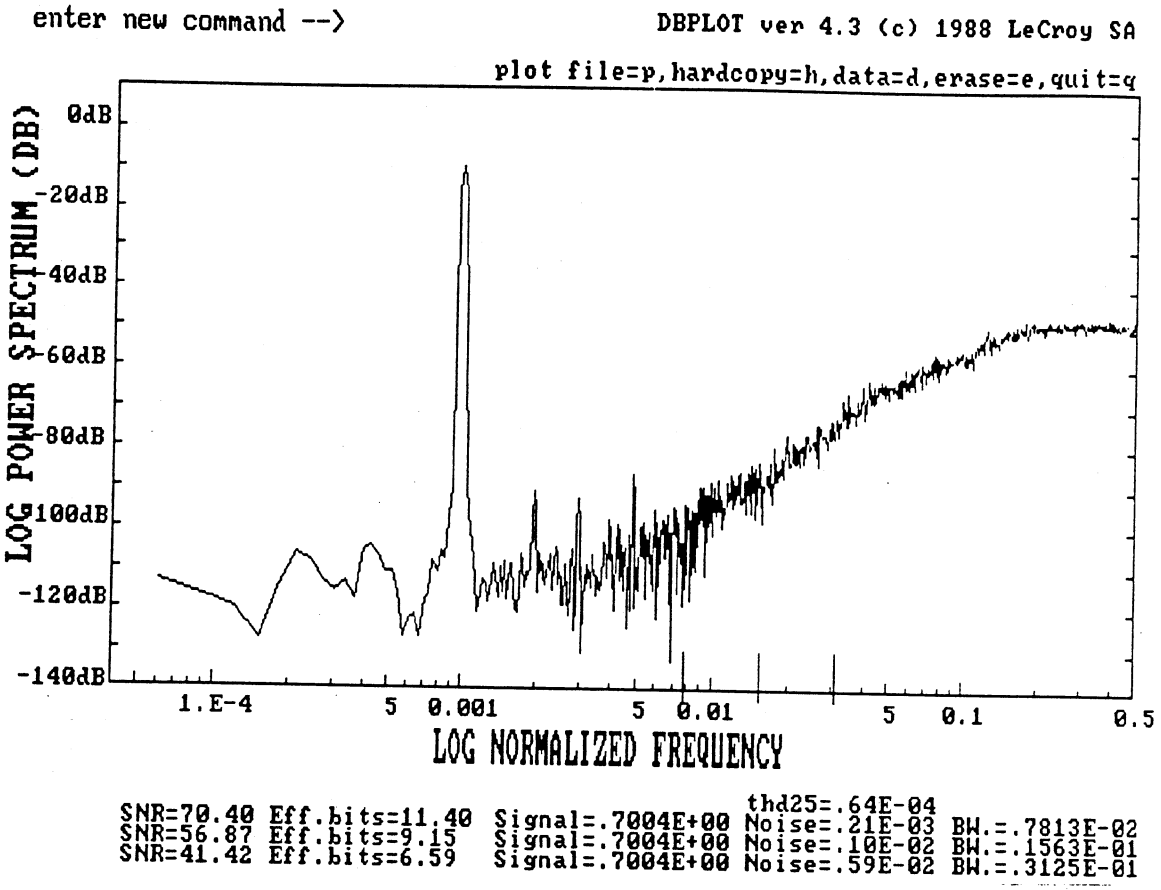
enter new command -->
DBPLOT ver 4.3 (c) 1988 LeCroy SA
plot file=p,hardcopy=h,data=d,erase=e,quit=q
LOG POWER SPECTRUM (DB)
0dB
-20dB
-40dB
-60dB
-80dB
-100dB
-120dB
-140dB
1.E-4
5
0.001
5
0.01
5
0.1
0.5
LOG NORMALIZED FREQUENCY
thd25=.64E-04
SNR=70.40 Eff.bits=11.40 Signal=.7004E+00 Noise=.21E-03 BW.=.7813E-02
SNR=56.87 Eff.bits=9.15 Signal=.7004E+00 Noise=.10E-02 BW.=.1563E-01
SNR=41.42 Eff.bits=6.59 Signal=.7004E+00 Noise=.59E-02 BW.=.3125E-01

Fig. 8

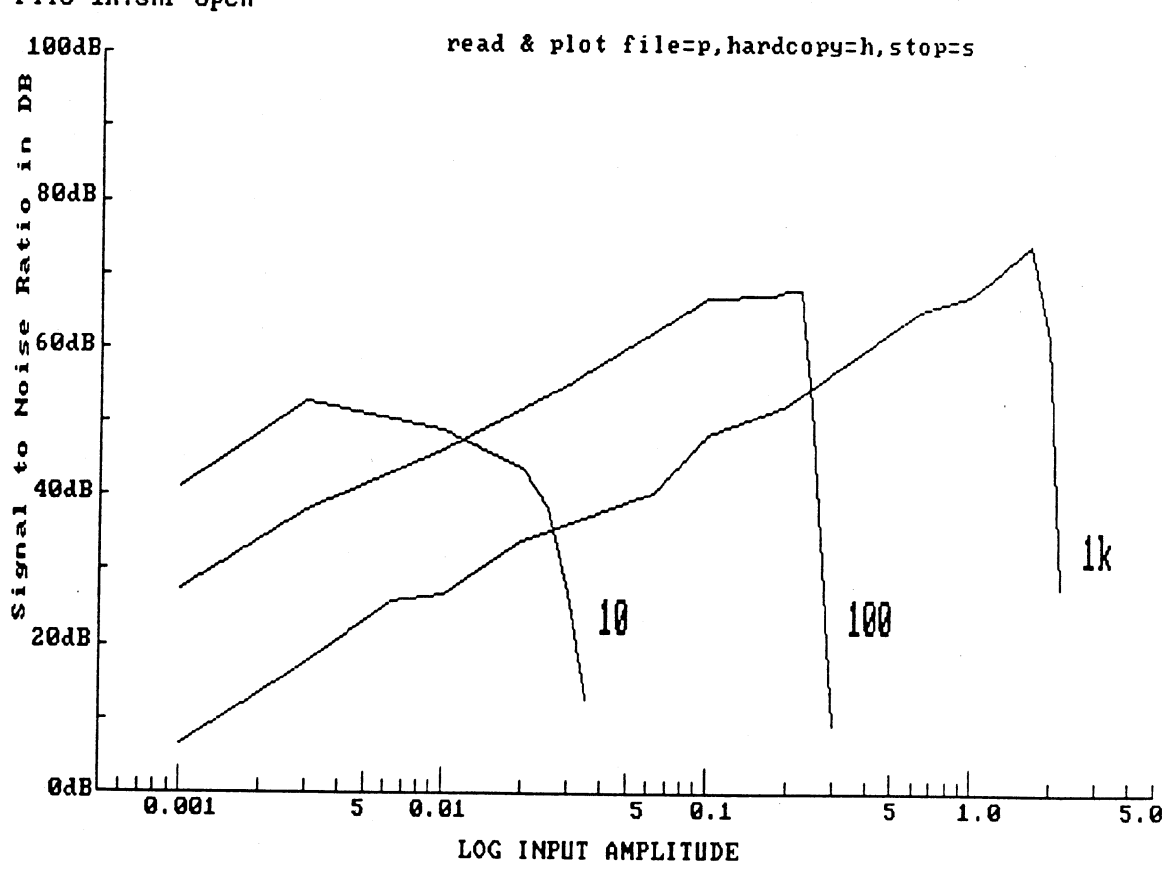
file 1k.snr open
read & plot file=p,hardcopy=h,stop=s
100dB
80dB
60dB
40dB
20dB
0dB
Signal to Noise Ratio in DB
10
100
1k
0.001
5
0.01
5
0.1
5
1.0
5.0
LOG INPUT AMPLITUDE

Fig. 9

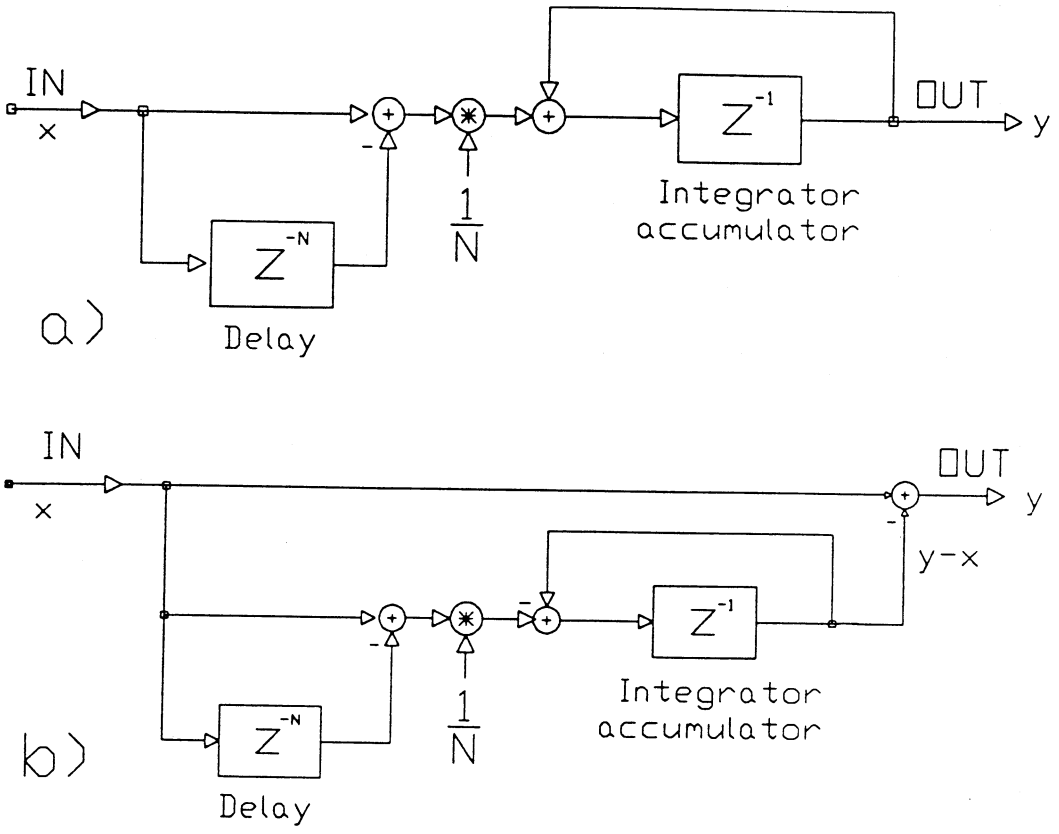

Fig. 10

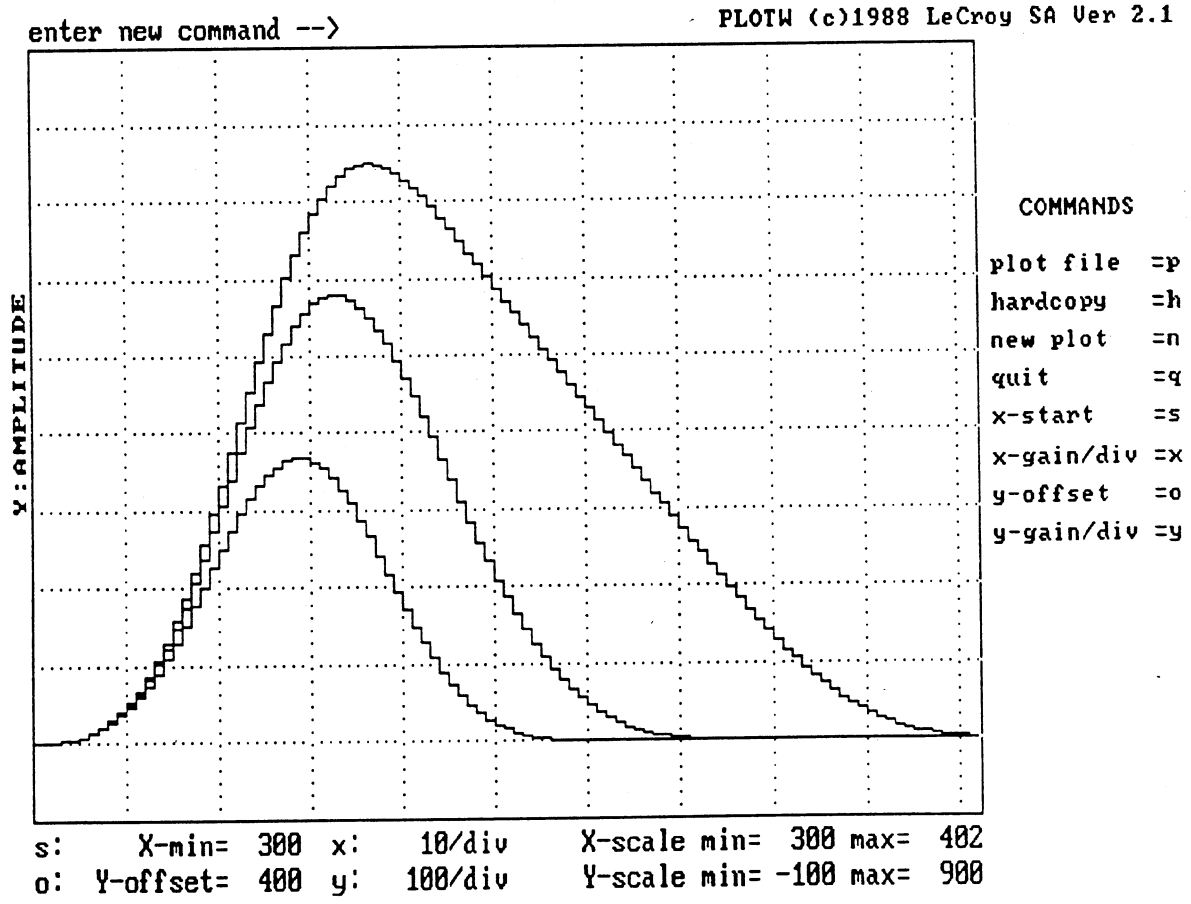

Fig. 11

Imprimé au CERN. Prix : 1.25 Fr.