Journal of Scientific Instruments (Journal of Physics E) 1968 Series 2 Volume 1

A wide range digitizer for direct coupled analogue signals

D N MacLennan and F H Wells
Electronics and Applied Physics Division, Atomic Energy Research Establishment, Harwell, Didcot, Berks.

MS *received 19 January 1967, in revised form 27 November 1967*

Abstract A serial analogue-to-digital converter is described which will accept currents in the range ± 1 mA applied to a virtual earth input, or alternatively voltages between ± 5 V applied to a high impedance input. An output pulse train is produced whose repetition frequency is proportional to the magnitude of the analogue input, the maximum frequency being 5 MHz corresponding to 1 mA or 5 V input. The conversion is linear to within $\pm 0{\cdot}2\%$, and the accuracy using a calibration curve is to $\pm 0{\cdot}1\%$ near full scale. The various factors which give rise to nonlinearity are discussed. The accuracy at a current input of 10 nA is to better than $\pm 10\%$, giving a dynamic range of at least 10^5 : 1 without the need for range changing. The corresponding figure for voltage inputs is much inferior to this, having a value in the worst case of only 3300 : 1.

1 Introduction

In many fields of applied science there is an increasing interest in the use of digital computers to process experimental data. The unaided computer, however, can only understand input which is presented in the form of a dimensionless number, while the data could be any kind of analogue information from a voltage level to a length or velocity. An 'analogue-to-digital converter' (ADC) is therefore required as part of the interface between the experiment and the computer. The digital output of the ADC could be transferred directly to the computer, or, for off-line operation, to some intermediate storage facility such as punched paper tape.

The successive approximation type of ADC (Gatti 1964) is the most widely used at the present time. The analogue signal is periodically sampled for a short time and the corresponding digital output is in parallel form. In the case of low frequency (down to d.c.) signals, such as one might get from a thermocouple or a photomultiplier, it is often more useful to know the average signal value over a period of time rather than the instantaneous value at intervals, for example, when high rejection of 'series mode' hum is necessary (MacLennan 1966). For these applications an 'integrating' ADC is required, and the instrument described here is of this type. The output is in serial form – a pulse train whose repetition frequency is proportional to the value of the analogue input, which is a voltage or current level in this case. The number of pulses produced in a given time is therefore proportional to the integral of the input level taken over that time, and so the mean input value is also known.

2 Current to frequency converter

The present instrument has been designed to digitize both current and voltage analogue signals. The basic principle of operation, however, is more suited to current (high source impedance) inputs. When it is required to digitize voltage (low source impedance) inputs, this difficulty is overcome by using a 'voltage-to-current converter' (as described in § 5) and the output from this converter can then be connected to the normal current input terminal. Thus the whole instrument is based on a current-to-frequency converter, and analogue inputs which cannot be applied directly to this circuit are first converted to a suitable form.

The current-to-frequency converter is shown as a block diagram in figure 1 (*a*), and some of the waveforms associated with this circuit are shown in figure 1(*b*); the particular waveforms given arise when the applied input current suddenly changes polarity. Although this does not often arise in practice, the example is useful to illustrate the bipolar operation of the circuit.

Consider first the time when I_{in} is positive. The d.c. amplifier and both charge feedback networks all present a very high impedance to the input, so I_{in} must flow into the tank capacitor C_{in}. The voltage across C_{in} therefore increases with time, at a rate proportional to I_{in}. This voltage is connected to the input of the d.c. amplifier and so the same waveform, greatly magnified, appears at the amplifier output. The discriminators (voltage) use tunnel diodes as the sensitive elements, as shown in figure 2, and are direct-coupled circuits designed to make use of the inherent backlash property of the tunnel diodes. Thus once the amplifier output voltage reaches a preset trigger level, the positive discriminator will be triggered, causing its output to change state. Being a direct-coupled circuit, the discriminator output voltage will of course remain in this new state until such time as the amplifier output voltage decreases below a reset level, when the discriminator will revert to its original state. The difference between the reset and trigger levels is the backlash of the circuit, which must be non-zero to prevent spurious oscillations or uncertain triggering.

The discriminator output controls a 5 MHz oscillator (figure 2); during the interval between the times of discriminator triggering and reset this oscillator will produce one or more pulses, pulses in the interval being separated by precisely 200 ns. The first pulse is coincident with the triggering of the discriminator, and the oscillator is switched off when the discriminator is reset. These pulses are fed to an output terminal and also to a charge feedback network. The operation of this part of the circuit is described in detail in the next section. To summarize, however, the application of

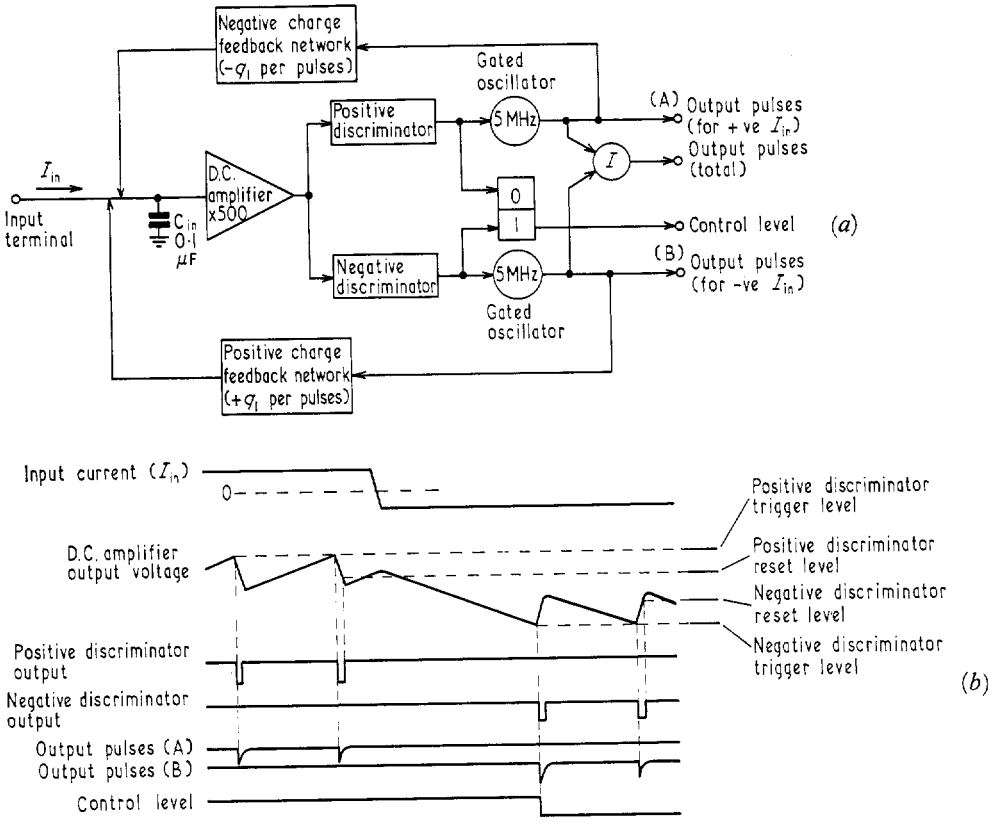

Figure 1 (*a*) Block diagram of current-to-frequency converter
(*b*) Idealized waveforms for input currents not approaching full scale

one oscillator pulse to the feedback network causes a defined quantity of charge opposite in polarity to I_{in} to be dumped on C_{in}. This process of charge transfer will be completed in about 50 ns, so the result is a fairly sudden reduction of the voltage on C_{in} and hence also of the d.c. amplifier output voltage. If the latter is now below the reset level of the discriminator, the feedback action of the circuit will cease with the oscillator switched off, only one pulse having been delivered to the output terminal. The whole cycle can then repeat itself. This will be the normal sequence of events for low to moderately high values of input current, as illustrated by figure 1(*b*). The capacitance value of C_{in} (a low inductance type) is chosen so that one pulse of charge feedback will normally take the discriminator from its trigger level to well below the next level, but not so far that the other discriminator becomes triggered.

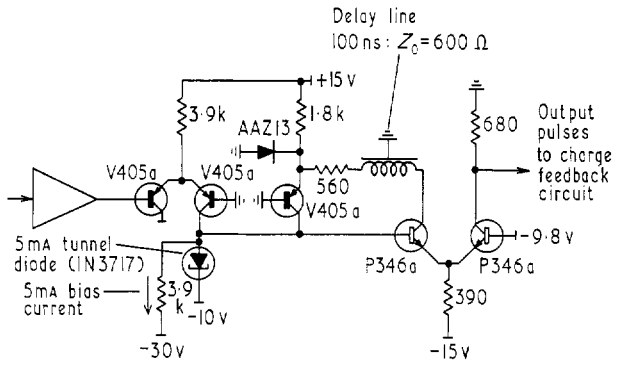

Figure 2 Positive discriminator and gated oscillator

In the case of input currents close to the full scale value of 1 mA the application of the first pulse to the charge feedback network may not be sufficient to reduce the amplifier output voltage below the reset level, and so the discriminator will remain triggered. A train of pulses at intervals of 200 ns will therefore be generated until sufficient charge has been dumped on C_{in} to allow the reset of the discriminator. For even higher values of I_{in} (greater than 1 mA) the amplifier saturates and the discriminator will not be reset at all. Thus the oscillator is permanently on and generates a continuous pulse train at 5 MHz.

When I_{in} has a negative value, a sequence similar to the above occurs. In this case, however, the amplifier output voltage will first decrease with time until it reaches the negative discriminator trigger level. Thus it is now the negative discriminator which is repetitively triggered and reset, the positive one remaining inactive. The negative discriminator is connected to another gated oscillator so that pulses corresponding to negative I_{in} appear at a separate output terminal. The two discriminator trigger levels are separated by about 2 v. This large separation is necessary to prevent spurious triggering due to the overshoots associated with the charge feedback action. The various levels which have been mentioned here are shown in figure 1(*b*).

The pulse trains corresponding to positive and negative I_{in} are combined by an OR gate to give a separate pulse output. Pulses from the latter thus represent a rectified version of I_{in}. An additional facility is the provision of a control level to indicate the input polarity. This is derived from a bistable which is controlled by the discriminators.

3 Charge feedback and linearity

The accuracy of the converter is determined principally by the constancy of q_1, the charge generated by the feedback network coincident with each output pulse. A variant of the well-known diode pump technique (Cooke-Yarborough and Pulsford 1951) is used to generate the charge pulses: the circuit of the positive charge feedback network is shown in figure 3; negative charge pulses are generated by a complementary circuit. The use of transistors to transfer the charge pulses to C_{in} means that the input of the converter always sees a high impedance, which is not the case when diodes are used. Moreover, the arrangement of figure 2 allows voltage excursions at the input of up to a few volts to take place whilst maintaining accurate operation of the converter. This situation could arise in conditions of transient overload input currents. The operation of the circuit is as follows. Initially J3 does not conduct, while J4 carries the small bleed current from R7. When a pulse is applied to the circuit input, the collector voltage of J1 rises by an amount $V_r + 2V_d$, where $2V_d$ is the combined forward voltage drop of diodes D1 and D2. A quantity of charge is then transferred to the input of the converter via C1 and the collector of J3; it is given approximately by

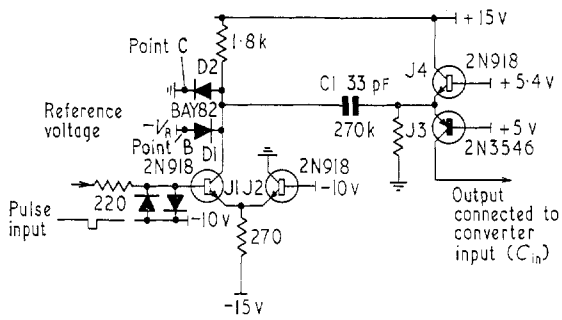

Figure 3 Positive charge feedback network

$$q_1 = \alpha C_1(V_r - 2V_d - 2V_{eb}) \qquad (1)$$

where $2V_{eb}$ is the combined forward voltage drops of the emitter-base junction diodes of transistors J3 and J4 and α is the grounded base current gain of J3. It will be seen that some degree of compensation for temperature effects has been achieved.

The principal aim of the design of this instrument has been to achieve the highest output frequency consistent with reasonable performance in other respects. It has been found that unacceptable nonlinearities arise if an output frequency of more than a few megahertz corresponding to the maximum input current of 1 mA is used. There are two known reasons for this. Firstly, charge storage effects in diodes D1 and D2. As the output frequency f increases (with increasing input current) the diodes have less time to dissipate their stored charge between pulses; the diode forward voltage drop on switching then becomes less, and so q_1 becomes dependent on f, decreasing as f increases. The transfer function, which is in its simplest form:

$$I_{in} = q_1 f \qquad (2)$$

therefore becomes nonlinear for large f. Initially IN 916 diodes were used, but q_1 was found to decrease by 1–2% between 1 and 5 MHz. The present circuit uses BAY 82 diodes which have a much lower stored charge, and these give a significant improvement in performance. Secondly, earth and power line ringing, owing to lead inductances and the very short pulse rise times which have to be used, also cause nonlinearities. As a change in q_1 of the order of 0·1% would be caused by only a few millivolts of disturbance at points B or C (figure 3) during switching, the importance of this effect will be appreciated.

Careful measurements of linearity were carried out on the prototype model, which has a maximum output frequency of 5 MHz, and the results are shown in figure 4. If the instrument is calibrated at an input current I_c to give an output frequency f_c the nonlinearity at any operating current of the same polarity as I_c is given by the equation

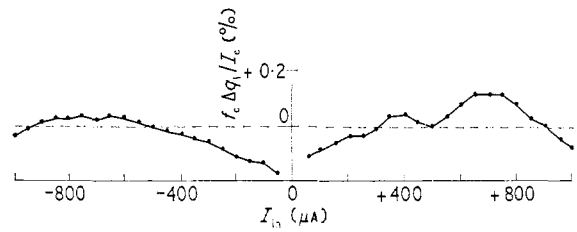

Figure 4 Conversion linearity plotted against input current $I_c = \pm 500\ \mu A$; $f_c = 2{\cdot}5$ MHz; measurement error, vertical axis, $< 0{\cdot}02\%$

$$\frac{f_c \Delta q_1}{I_c} = \frac{f_c}{I_c}\frac{I_{in}}{f} - 1 \qquad (3)$$

where Δq_1 is the change in q_1 relative to its calibrated value; f was measured simply by counting the output pulses in a scaler for a precise period of time, determined by a crystal oscillator. To resolve the fine structure of the nonlinearity curve, I_{in} must also be measured digitally, and this was done by connecting a digital voltmeter across a 1000 Ω high stability resistor in series with the converter input.

An examination of the curves in figure 4 shows that q_1 does in fact tend to decrease with f at high output frequencies. Fluctuations which might be due to the second effect discussed above are also evident, although these exceed the measurement error only in the case of positive input currents.

Nonlinearities which are due to the high frequency limitations of the circuits would be expected to be negligible at small values of I_{in}. However, q_1 appears to have a roughly linear dependence on f in this region which persists even at the lowest frequencies. The most likely cause of this variation is thermal effects, which are discussed in another section.

4 Dynamic range (current input)

The lower end of the useful operating range which sets a maximum limit to the dynamic range is determined by drifting of the total leakage current which flows into C_{in} when the input is open circuit. Neglecting leakage resistances this current will arise from the four circuit components which are attached to C_{in}. The d.c. amplifier has a differential field effect transistor input stage, and the input current to this transistor will drift by less than 0·1 nA $degC^{-1}$ and can be neglected. Zero setting is accomplished by controlling a small bleed current into C_{in} through a 100 MΩ resistor, but drifts due to this are also small. The most important effect is due to the collector leakage currents of J3 and its complementary, which oppose each other. I_{CO} for these transistors has a maximum value of 10 nA, giving a drift in the worst case of not more than 1 nA $degC^{-1}$. A typical figure would, of course, be much less than this (Owens and Perry 1965) and indeed the prototype model was found to have a measured input-referred current drift of only 50 pA $degC^{-1}$.

5 Voltage to frequency converter

The voltage from a low impedance source could be digitized

by connecting a series resistor between the source and a current-to-frequency converter (de Sa 1965). In the present case, a 5 kΩ resistor would give a voltage range of ±5 v. A major disadvantage of this method, however, is the existence of a 'dead zone' about 0 v. This arises because the voltage on C_{in} must be moved from its quiescent value by a finite amount before a discriminator is triggered. As the d.c. amplifier (figure 1) has a gain of 500 and the discriminator thresholds are separated by about 2 v no output pulses would be produced for voltage inputs in the range ±2 mv, giving a severely limited dynamic range. A significant improvement in this respect is achieved by using the arrangement shown in figure 5. R6 is adjusted so that the loop gain L, from one end of R4 to the other, is close to unity. L is given by the equation

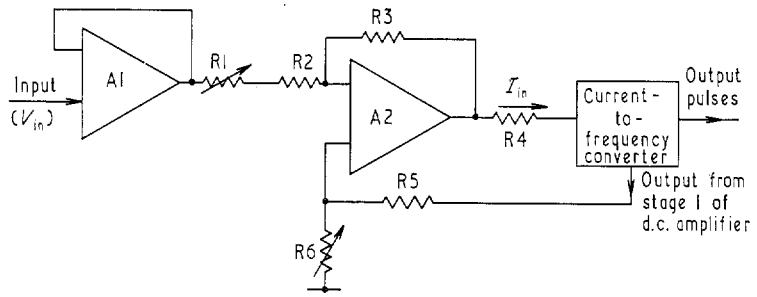

Figure 5 Voltage-to-frequency converter

$$L = a_1 \frac{R_6}{R_5 + R_6} \frac{R_1 + R_2 + R_3}{R_1 + R_2} \tag{4}$$

where a_1 is the gain of stage one of the d.c. amplifier. A high input impedance is achieved by the use of buffer amplifier A1, which also serves to make L independent of the resistance of the voltage source. If the voltage on C_{in} is V_i, the current through R4 is given by

$$I_{in} = -V_{in} \frac{R_3}{R_4(R_1 + R_2)} + \frac{L - 1}{R_4} V_i. \tag{5}$$

Thus, with $L = 1$, I_{in} is independent of V_i and the dead zone is entirely eliminated. In practice, however, L must be made slightly less than unity to prevent possible instability, but a value of at least 0·99 is quite feasible, giving a maximum dead zone of only ±20 μv.

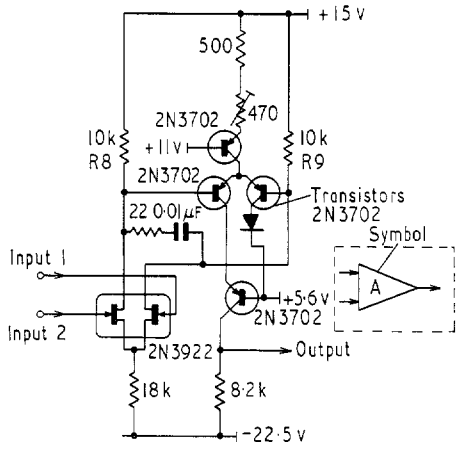

Figure 6 Differential amplifier circuit

Calibration of the circuit is effected by varying R1, which is set so that a voltage input of ±5 v corresponds to a current through R4 of ∓1 mA.

6 Dynamic range (voltage input)

The lower end of the useful operating (voltage) range is determined by drifting of the input offset voltage γ_0 which is the value of V_{in} required to give zero output frequency. This arises from the combined effect of offset voltage drifts in the three amplifiers associated with the voltage input, namely A1, A2 and stage one of the current-to-frequency converter d.c. amplifier. The same basic amplifier circuit, shown in figure 6, is used for each of these. The differential inputs are connected to the gates of a dual field effect transistor, type 2N3922. This device has a differential gate voltage drift of less than 25 μv degc^{-1} when used with the correct drain current. The output is taken from the collector of grounded base stage J7. Any required gain which is less than the open loop figure can be obtained from this circuit by connecting appropriate feedback resistors between the output, one of the inputs and ground.

The input to the field effect transistor that is necessary to achieve high input impedance gives a much lower stage gain than would be obtained by using bipolar transistors. For this reason input referred drifts due to the rest of the circuit are not negligible. To reduce these to a satisfactory level, high stability resistors (±25 p.p.m. per degc) have been used for R8 and R10. Also transistors J5 and J6 are matched to have a common emitter current gain ratio of 1 ± 0·1, and R10 is adjusted (with feedback applied) so that J5 and J6 collector currents differ by less than 5%. Thus temperature dependent changes in the base currents of the two transistors track closely.

If the voltage offset (referred to their own inputs) of the three amplifiers referred to above are γ_{01}, γ_{02} and γ_{03} respectively, the total voltage offset drift referred to the V_{in} terminal will be given by

$$\frac{\delta\gamma_0}{\delta T} = \frac{\delta\gamma_{01}}{\delta T} + \frac{\delta\gamma_{02}}{\delta T} \frac{R_1 - R_2 + R_3}{R_3} - \frac{\delta\gamma_{03}}{\delta T} \frac{R_1 + R_2}{R_3} \tag{6}$$

where T is the temperature. Each of the three terms in this equation will in turn be related to drift contributions from the individual components of the amplifier circuit of figure 6, which would in general be expected to vary randomly in magnitude and sign from one amplifier to the next. A maximum value for $\delta\gamma_0/\delta T$ could therefore be arrived at by summing the 'worst case' contributions from these components regardless of sign. In view of the large number of components involved, however, such a figure would be unrealistic. A better method is to carry out a statistical analysis of these individual effects, bearing in mind their distribution within their maximum and minimum limits. The latter will

usually have some assumed idealized form depending on the manufacturing technique used. Calculations on this basis indicate that the distribution of $\delta\gamma_0/\delta T$ should be approximately Gaussian in form, with three-sigma limits at

$$\pm 150\ \mu\text{v degc}^{-1}.$$

This is in good agreement with the measured drift of the prototype model, which is $+70\ \mu$v degc^{-1}.

On the reasonable hypothesis that the derived figure of 150 μv degc^{-1} represents the worst case, the dynamic range would not be less than 3300 : 1, assuming the circuit temperature to be controlled to within ± 1 degc and a measurement accuracy to $\pm 10\%$ at the bottom end of the range.

7 Thermal effects

The effect of ambient temperature variation has been reduced as far as possible by circuit design compensations but the performance would be much improved by stabilizing the temperature to within ± 1 degc. Some attempt at this stabilization by using oven control was only partially successful owing to thermal gradients in the oven, and this oven has not been used.

A further difficulty is that since most of the circuit is electrically direct-coupled, changes in the input signal level will give rise to corresponding power dissipation changes in both active and passive circuit components, which will in turn cause temperature changes at various critical points in the circuit.

Referring to figure 3 and equation (1), there are many components which could affect the conversion factor by this process. Assuming lead conducted heat to be the controlling factor, diodes D1 and D2 would be expected to cause the greatest effect because of the relatively high power dissipation in the components connected to the collector of J1. J1 is normally ON, and is only switched OFF for a short time coincident with each output pulse. To a first approximation, therefore, the heat to be dissipated by D1 and D2 should decrease linearly with increasing input magnitude. The resultant change in the diode forward voltage drop ($+2$ mv degc^{-1} fall) is in the right direction to explain the residual nonlinearity found at low input currents that was mentioned in § 3. The thermal origin of residual nonlinearity is also suggested by the response of the converter to a step function input (figure 7). Following the application to the

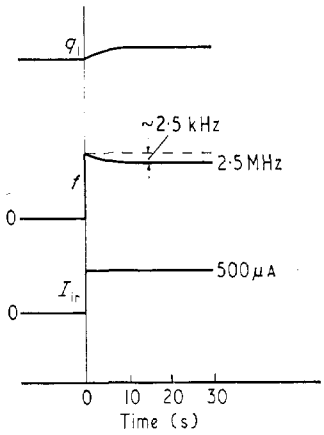

Figure 7 Response to a step function current input

input of a current step from 0 to 500 μA (either polarity) the output frequency overshoots its final value by about 0·1%. This excess frequency decays with a time constant of several seconds. The time dependence of q_1 can be deduced from this information by using equation (2); it is shown in the top curve of figure 7. As the measured points in figure 4 are sufficiently separated in time for transient effects to have disappeared, the excellent correlation between the transient frequency overshoots and the observed nonlinearity in the current range $\pm 500\ \mu$A suggests that the response to pulsed current inputs in this range would be quite linear, provided of course that the mark-to-space ratio of the pulses was always much less than unity.

Various forms of heat sinking were tried in an attempt to reduce the effect of dissipation changes, but without much success. Another approach is to reduce the current and voltage levels in critical components as much as possible. Unfortunately, the high frequency performance of the transistors used suffers when the operating collector current and voltage levels are reduced below the recommended values. Thus any reduction in thermal effects would probably be offset by an increase in frequency dependent nonlinearities, unless a lower maximum operating frequency were accepted.

8 Performance

When used as a current-to-frequency converter, the 2192-2 ADC produces an output pulse train whose frequency is proportional to the input current. The maximum output frequency is 5 MHz, corresponding to an input current of 1 mA, and the conversion is linear to within $\pm 0{\cdot}2\%$. For the highest working accuracy a calibration curve should be used to allow for the nonlinearities inherent in the system; then the overall accuracy would be to better than $\pm 0{\cdot}1\%$ at high values of input current. The accuracy for lower currents is reduced by the effect of leakage current drifts at the input, but should be to considerably better than $\pm 10\%$ at 10 nA. Alternatively the ADC can be used as a voltage-to-frequency converter, simply by making an external connection on the front panel and applying the signal voltage to a second input terminal. A negligible amount of nonlinearity is introduced by the additional circuitry required in this mode, and so the linearity of response is essentially the same as for current inputs ($\pm 0{\cdot}2\%$). Using a calibration curve the conversion accuracy will be to within $\pm 0{\cdot}1\%$ at 5 v, and $\pm 10\%$ at 0·15 mv.

The dynamic range for voltage inputs (3300 : 1) is much inferior to that achieved with current inputs (10^5 : 1), being nearly two orders of magnitude less. This is due to offset voltage drifts in the linear amplifiers which are used in the voltage mode. A considerable improvement in performance would be expected if the direct-coupled amplifiers of the present design were replaced by chopper type amplifiers, which can have offset voltage drifts of the order of 1 μv degc^{-1} or better. This would, however, involve a considerably more complex circuit than the present one, and the cost of the instrument would be greatly increased.

Acknowledgments

The development of the instrument described here was carried out by E.M.I. Ltd. in conjunction with the U.K. Atomic Energy Authority. We should like to thank our colleagues, Mr G Prior and Mr N White of this firm, for their part in bringing this project to a successful conclusion.

References

Cooke-Yarborough E H and Pulsford E W 1951 *Proc. Inst. Elect. Engrs* **98** 191–5

de Sa A 1965 *J. Sci. Instrum.* **42** 265–7

Gatti E 1964 *Proc. Conf. on Automatic Acquisition and Reduction of Nuclear Data* (Karlsruhe: Gesellschaft fur Kernforshung). p. 259

MacLennan D N 1966 *Atomic Energy Research Establishment Report No.* R 5306

Owens A R and Perry M A 1965 *Proc. Inst. Elect. Engrs* **112** 2005–12