Project 1523

Digital readout electronic scales

Part 1

Ian Thomas

FOR MANY YEARS I've been toying with the idea of building an electronic weighing scale but up until now I've always put it in the "too hard basket". However, after reading an article in Wireless World by John L. Linsley Hood (*Strain-Gauge Weighing Scale,* Wireless World October 1983) I decided that the time was right. It also seemed a nice idea to have a scale that had multiple ranges to be used for weights from less than a gram to 5 kg and so the specification for my scale began to take shape. As there are a multiplicity of 3½-digit DVM chips available the scale naturally would have ranges of 200 gm, 2 kg and (hopefully) 20 kg with resolutions of 0.1 gm, 1 gm and 10 gm respectively (although the 20 kg proved too much and I had to settle for 5 kg).

All electronic scales consist of a transducer to convert the gravitational force produced by the mass of the item being weighed, electronics to amplify and condition the transducer output and a digital display system to show the result. While the electronics presented no insoluble problems the transducer to convert force to an electrical signal was a different story. While there are many ways of constructing such transducers most require access to sophisticated tools and technology which most of us (myself included) don't have.

The fact that this project has to be buildable by the home constructor eliminated most options such as linear variable differential transformers (LVDTs) or linear potentiometers, as even if one could be made at home you would have a snowball's chance in hell of making it linear enough for this application (a 0.1% linearity LVDT costs more than you or I are prepared to pay and anyway it would take all the fun out of it to buy the heart of the scale!). All of these factors forced me to the conclusion that the right way to construct the transducer was to use some sort of spring which deflects under the load and measure the deflection with a strain gauge.

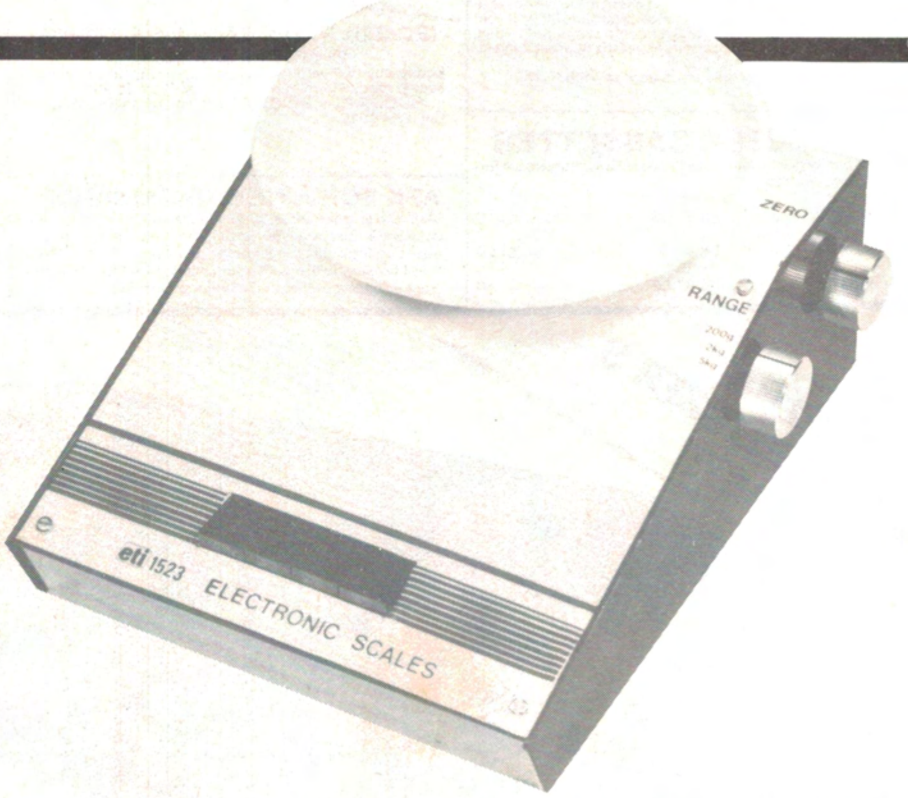

Clean lines. I housed the project in a Bimbox which gives clean lines and an ergonomically satisfactory layout.

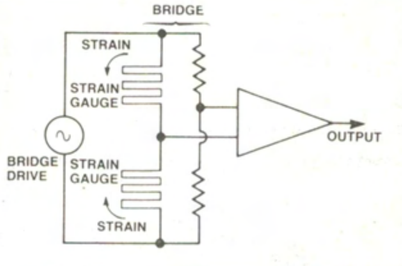

Figure 1. Strain gauge 'bridge' system. Each strain sensor experiences the opposite strain, unbalancing the bridge. The output is then amplified.

The strain gauge

For those of you who have not yet run across the term "strain" in the mechanical sense it is defined as the elongation of a member under load divided by the length of the member (mathematically dl/l) and is usually expressed as a percentage. For most springs the ratio of force applied to the strain produced is very linear (the spring is said to be "linear elastic"). A strain gauge works on the very simple principle that if you pull hard on a piece of conducting material then it gets slightly longer and slightly thinner which causes its resistance to increase. Exactly the opposite occurs if you compress it. Because of the lengthening *and* thinning a 1% strain will cause a 2% resistance increase in the material and the whole effect will easily give the linearity we need.

Given that the scale was to be of the strain gauge type the next problem was to decide on the mechanical structure of the spring strain gauge combination. Also, as the strains to be measured are *very* small and there are other effects that change the resistance of conductors (like mainly, temperature) it is infinitely desirable to use strain gauges in a bridge structure where one gauge is compressed and another is stretched by the same load. Other extraneous effects such as temperature should (hopefully) affect both gauges equally and cancel (see Figure 1). This gives the second requirement for the spring structure; it must allow two gauges to be mounted in juxtaposition so they experience strains of opposite sign.

The simplest type of spring I could think of that fulfilled all of these requirements is the simple cantilever. A cantilever is just a

Employing a unique sensing technique, with a strain gauge printed on the pc board, this project avoids the necessity of using difficult to get strain gauge sensors, linear pots, LVDTs etc. It has reasonable precision, four-digit readout, and three ranges of 200 gm, 2 kg and 5 kg full-scale.

bar held horizontally and rigidly clamped at one end. When a load is applied to the other end the top surface of the bar is stretched and the bottom surface is compressed equally; which is exactly the result we want. Therefore, a cantilever or combination thereof seemed to be the right way to go. The next problem to be addressed was what sort of strain gauge to use. Commercially available (and quite cheap) strain gauges consist of foils of fancy alloys bonded to plastic film which are glued to the test piece. The foils are etched in meander line patterns so the long runs of the meander are in the direction of the strain to be measured and strains at 90° to the meander produce (almost) no effect. A possible answer would have been to simply purchase some of these strain gauges and stick them to some sort of spring but I come from a long line of tightwads and wanted a cheaper answer.

The printed circuit strain gauge bridge

I've noticed many times just how surprisingly strong and springy normal epoxy-glass printed circuit board material is and it seemed to me that it would make ideal material for the spring cantilever(s) of the scale. The next obvious thought was not to glue foil strain gauges to the surface but to use the foil that was already there, namely the copper. A meander line structure could be etched in the copper cladding to produce (free!) strain gauges as needed. A further advantage of this structure would be that the foils on both sides of the laminate are thermally in close contact and should track each other. A quick test board was made with only strain gauges on it and lo and behold, it worked! I etched the same pattern on both sides of a 20 mm x 50 mm piece of board as you can see on one arm of the final artwork and used each side as one arm of a bridge. With suitable excitation of the bridge (a 5 kHz square wave with as much power as I could use without burning things up) quite useable outputs were obtained when the board was bent. The only problem was that the two meander line patterns had a very low resistance — about two ohms each.

Inside. Showing general layout and construction. The cotton wool hiding the strain gauge transducer provides thermal insulation to reduce temperature effects.

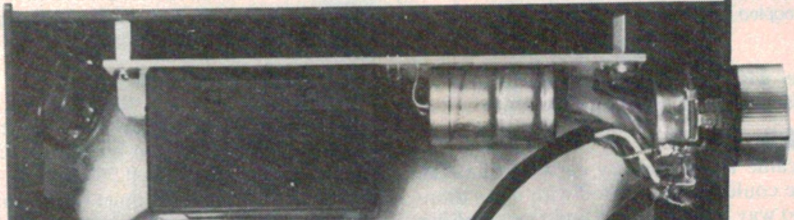

The low resistance presented a problem for the following reasons. If you have a bridge that is perfectly balanced and apply an exciting voltage to it you get nothing out. If you then unbalance it by increasing one arm resistance by 1% and decreasing the opposite arm's resistance by 1% then the output is ½% of the exciting voltage regardless of the actual value of the resistors. Therefore, the exciting voltage should be as big as possible to improve the signal-to-noise ratio of the transducer. The only limit to the magnitude of the exciting voltage is how much power you can drop in the bridge arms. Very low arm resistances mean very high power dissipation in the bridge to get a good signal-to-noise ratio for the transducer; hence the problem.

As I mentioned earlier, temperature effects also change the arm resistances and the very last thing wanted is to have the arms glowing a dull cherry red when you're trying to sniff out microvolt signals!

It's interesting to put a few numbers to this problem to illustrate it. If the arm resistances are one ohm and excited with one volt then they dissipate one watt, which makes them quite hot. If our maximum load of 20 kg gives a 5% resistance change then the bridge output will be 2.5% of one volt, or 25 mV. However, our desired resolution is 0.1 gram or 0.000005 x 25 mV or 125 nanovolts, which is stretching the friendship a bit. To further compound the problem it is necessary to increase the arm resistance four times to double the exciting voltage as power is proportional to E^2. All this says is that, so far as arm resistance is concerned more is better, and even more is better still!

When I sat down and calculated the absolute maximum track length I could cram onto a suitable size cantilever arm the answer came out only just possible, but achievable. A further arm resistance ▶

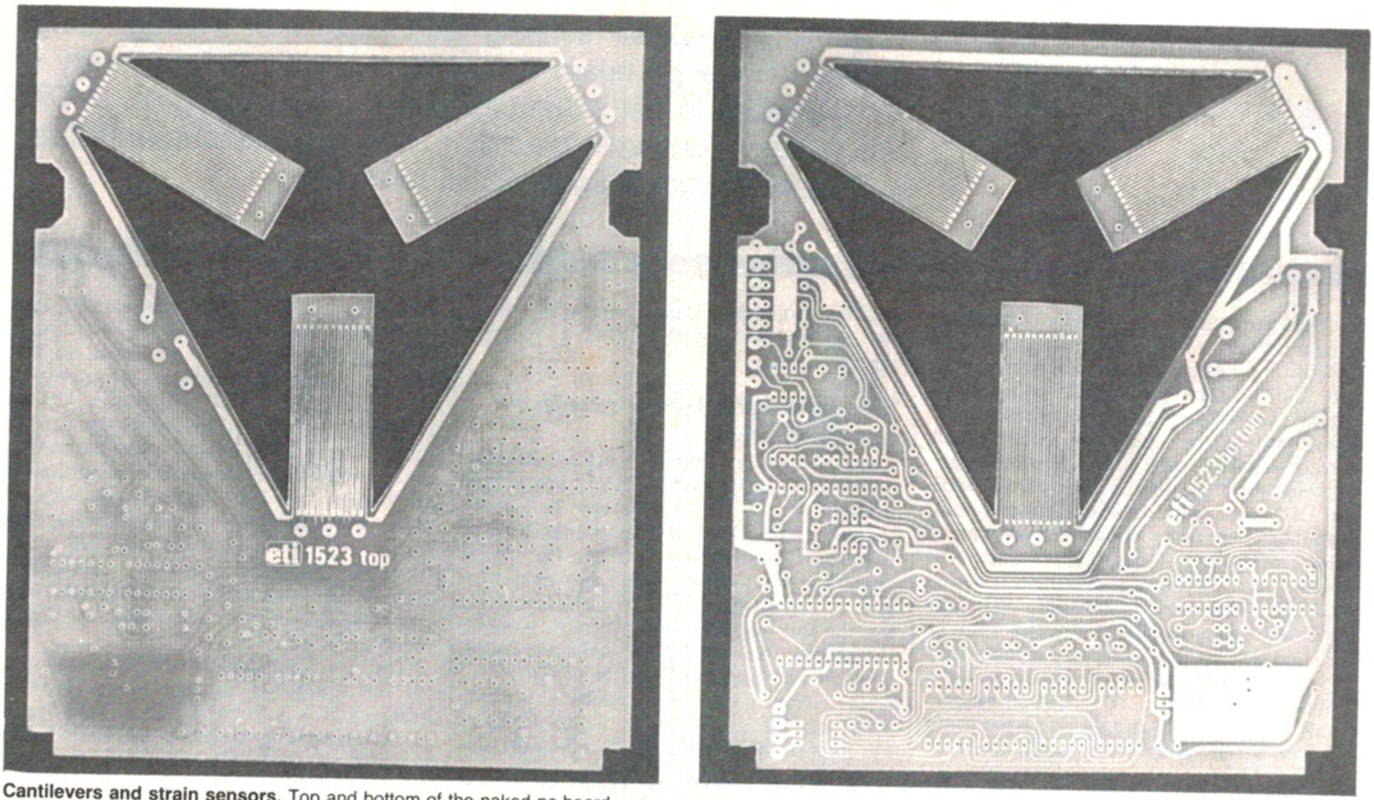

Cantilevers and strain sensors. Top and bottom of the naked pc board showing the three cantilever arms and the strain gauge meander tracks. Careful board layout has obviated problems with noise being coupled into the sensitive bridge amplifiers.

increase could be achieved by using the thinnest copper laminate available ("½ oz" or 18 μm thick copper). By using 0.5 mm track widths and 0.38 mm spacing a 50 mm x 20 mm area of meander line would have a resistance of about three to four ohms — just usable. In the final design, where there are three separate cantilevers, the total arm resistance worked out at about 10 ohms per side which was (sort of) all right but then created the next problem; how to excite the bridge without wasting power (power-is-heat-is-trouble in anything this sensitive).

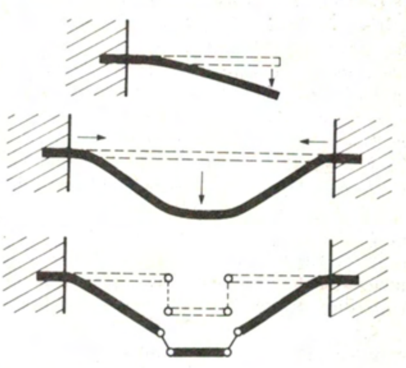

Figure 2. Evolution of the strain transducer: A single cantilever (top) to a dual cantilever (middle) to a multiple cantilever system with pivotted load support.

Bridge drive

It became apparent very early on that the bridge could not be excited with dc as there was no way I was going to build an amplifier with input offset voltages of 100 nV or less. Life is too short as it is. With a total bridge arm resistance of about 20 ohms a casual poke at a calculator reveals that we want an exciting voltage of three to four volts before things start to become awkwardly warm. Given that the power supply gives 15 volts (actually ±8 for reasons that will be discussed later) a method was needed to derive a three or four volt signal with heaps of drive capability without wasting three times as much power in the drive circuitry by straight regulating down. The obvious answer was to use ac drive and use a transformer. This meant that the drive voltage could be varied simply by changing the secondary turns ratio and would give the desired efficiency.

Given that the excitation was to be provided from a transformer the next decision was what frequency to use. The upper limit was set by the bandwidth of the amplifiers following the strain gauge bridge and the lower limit was set by the size of the transformer (or, more accurately, the number of primary turns — I personally find winding toroidal transformers pure, unremitting boredom!). The two limits set the drive frequency at between two and 10 kHz and I finally settled on about 5 kHz.

The load transducer

The heart of the scales is the transducer that converts the applied load to an electrical signal for processing and measurement. It must be linear to at least 0.5% for a resolution of 1 part in 2000; It must be capable of handling extreme overloads without damage (in case someone drops a brick on the weighing pan) and for our application it must be dirt cheap. All in all a challenging specification. As has already been discussed these needs didn't leave a lot of choice except a strain gauge type of transducer and a cantilever spring so the gauges could be mounted in opposite load positions to make a bridge circuit possible. Given that this was what was wanted a mechanical configuration that would tolerate all manner of abuse had to be contrived (I realise that you, dear reader, will be careful, but others may not!). The structure also had to provide support for the weighing pan that was reasonably rigid and, if possible, not consist of any sliding members or messy mechanics that would give hysteresis or tend to stick. One simple cantilever would do this except that when a heavy load was applied the weighing pan would droop sideways — not so good.

I then contrived the idea of having three or four separate cantilever beams in a radial pattern that were rigidly mounted at their outer edges and rigidly attached to the weighing pan in the centre (see Figure 2). This would certainly meet the mechanical ▶

requirements but had the problems that it was too strong and would not deflect enough (remember that strain gauges measure deflection) and would give heavy side loads to the board supports. Also it would require that the strain gauges be separated into outer and inner halves and carefully connected up so the right gauges were in the right arm of the bridge. The connection problem in itself wasn't major but it made making the artwork for the board messy.

The final answer that seemed to solve all problems was to separate out the three inner ends of the cantilevers and fix them with pivots to the weighing pan support. Since there were three points of attachment to the weighing pan there should be enough sideways support to stop the pan tilting; and also, if the pan did tilt a bit because of an off-centre load, the strain gauges are summed so the overall load measured would be correct. The only problem with this structure was that there are mechanical pivots carrying the load. A quick trip to my friendly local hardware store told what was available here (to be candid, not very much — whatever happened to the shops where you could buy just about anything!). However, I could get ¾″ hinges that could possibly be made to do the job. When they came out of the packet they were far too stiff and would have probably given the scale some hysteresis but after oiling and working them they freed up enough to be tried. As the load ends of the cantilever beams move down and sideways under load two hinges had to be used per beam but the bank could stand the expense.

The outer ends of the three beams had to be rigidly attached to some form of base plate and once again I didn't want to get involved with complicated mechanics. Since almost everyone uses tapped spacers and the steel ones are *very* strong this seemed the easiest way to go. I mounted the board on a solid aluminium base with three spacers instead of one at the clamped ends of the cantilever beams. I suspect the resulting structure would survive having a truck driven over it; it certainly was rigid enough. The base plate itself is just a 160 mm square of 2 mm thick aluminium with assorted holes and notches cut in it — no problem.

When I was starting to put this project together I was trying to keep all unnecessary weight off the centre and weighing pan support so I had a threaded bar of aluminium made up to be attached to the centre of the three beams and support the weighing pan but I really think this is unnecessary. A simple ¼″ bolt will do just fine (at least they're easy to get). You need two nuts to attach the lower end to the centre of the beams and I just Araldited a third nut to the bottom of the weighing pan to screw it onto the shaft.

The shaft that holds the weighing pan also very nicely provides an end stop for the travel of the scale. If the beams are deflected too far the bottom of the shaft hits the aluminium base plate and prevents anything being broken. I haven't tried dropping a brick on it yet but I suspect it would survive. The whole structure seemed to meet all requirements very nicely indeed.

Reading the strain

As I've already mentioned, we have to amplify *very* low level signals in the scales so the whole electronics design has to be oriented around low noise, accurate performance. To this end it is essential that the earth for ac signals be the earth for everything; hence the split rail power supply. As a general rule "almost earths" or "not quite earths", formed by resistive dividers, are a recipe for trouble in low level systems and should be avoided like the plague. For the cost of one more filter capacitor and a transformer with a centre tap the problem can easily be sidestepped. Since the circuitry has some CMOS, the split rails were kept to ±8 volts which is also quite adequate to power the op-amps. Another very nice thing to do in any design is eliminate the need for regulated power supplies. In this case the exciting voltage is directly related to the supply so the processed output voltage from the bridge would also be proportional to the supply.

However, the cheaper digital voltmeter chips that don't have an internal voltage reference inherently give a digital output that is proportional to the ratio of the input to reference voltage. This means that if the reference voltage is proportional to the supply voltage then the DVM output would be proportional only to the strain — exactly what is wanted! In fact I took this one step further and derived the reference voltage in exactly the same way as the output voltage is processed, in order that the generated reference voltage follow exactly all variations of both the positive and negative rails.

It proved to be rather fortunate that this was done as the digital LED display causes the unregulated positive rail in the final scales to move around quite markedly. The reference and output voltages being held exactly proportional to the supply completely cancel this out and the display shows almost no tendency to affect input levels.

The drive voltage for the transformer was supplied from between the two rails rather than between one rail and earth. having generated a real "earth" earth, the *last* thing to do is knowingly dump very noisy currents onto it so different parts of the same earth track are at different voltages due to IR drops along the track. The most obvious way of generating the excitation voltage is to use a simple saturating core inverter and don't bother with the output diodes.

I've learned from past experience that when saturating core inverter cores saturate

▶

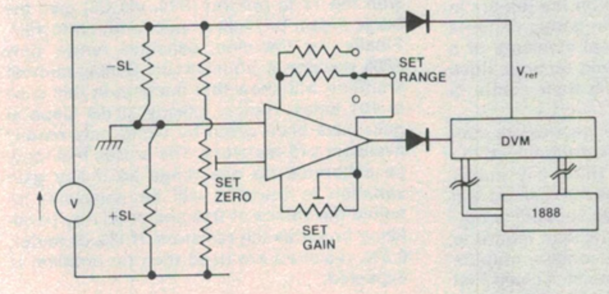

Figure 3. Block diagram of the project.

Load support. View of the completed transducer showing the pan support bolt and hinge pivot system. Note the cantilever end supports at top left.

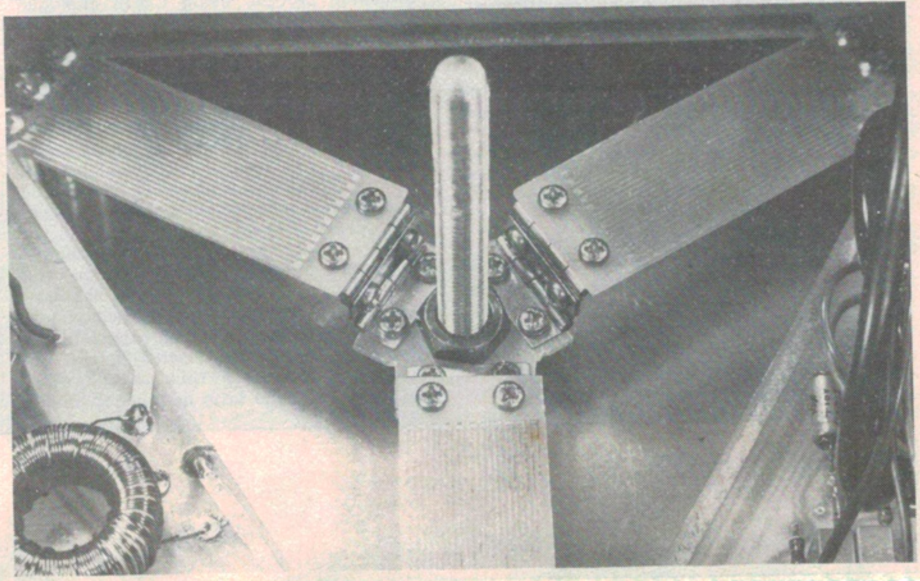

HOW IT WORKS — ETI-1523

The project can be divided into six separate sections:

1. The power supply;
2. The strain gauge transducer and its associated mechanics;
3. The strain gauge oscillator and bridge drive circuit;
4. The strain gauge output amplifier;
5. The synchronous switch for dc recovery of the amplified ac signal and the associated output dc differential amplifiers;
6. The digital voltmeter and display.

Each part of the circuit will be discussed in turn.

POWER SUPPLY

The power supply is a simple centre-tapped 6 V secondary transformer followed by a full wave rectifier comprising diodes D1 to D4. As the transformer centre tap is connected to earth the two rectified dc outputs take up voltages of ±8 volts for use in the rest of the system. Large filter capacitors, C18 and C19, smooth the power supply sufficiently for use in the project without further filtering or regulation.

THE STRAIN GAUGE

The strain gauge transducer is the heart of the scale and its accuracy and linearity set the performance of the whole instrument. It consists of two groups of three 'meander line' pattern strain gauges etched on 'half-ounce' copper laminate. A meander line is formed on each side of three cantilever beams cut out in the laminate which are arranged in a radial pattern with their outer ends rigidly clamped. The inner ends of the three cantilever beams are free to move under the applied load and are joined by a hinge structure that supports the weighing pan.

When a load is applied all the strain gauges on the bottom side of the beams are compressed and all the gauges on the top are in tension. Thus the two strain gauge patterns on either side form the ideal elements of a bridge structure. The patterns on both sides are connected in series and their centre is connected to ground.

The two ends of the strain gauges are connected to a floating secondary winding of the bridge drive transformer. Thus, any imbalance in the two strain gauge resistances will cause the voltage on the two secondary outputs to be slightly assymetric with regard to ground. These same two secondary outputs are connected to a series resistor chain that has a centre tap that is adjustable. The centre tap is the output to the electronics and the adjustment is to enable any zero offsets to be nulled out.

The inner ends of the three beams are attached to each other via a system of hinges that enable the ends to move freely away from each other when the beams are flexed but still provides mechanical support to the weighing pan. In order to minimise thermal effects the three strain gauge beams are wrapped in a thick layer of cotton wool. This limits heat loss to the ends of the beams and ensures that there are no rapid temperature changes between two strain gauges (see article for details of construction).

OSCILLATOR AND BRIDGE DRIVE

The drive signal for the strain gauges is derived from a simple CMOS oscillator made up from two NOR gates in IC1. The frequency of oscillation is determined by R1 and C2. The output of the oscillator, pin3 of IC1, is divided by one half of the type D flip-flop IC2. This ensures that the output to be fed to the toroid drivers (Q1, Q2) is completely symmetric. The output pins (12 and 13) of the divider are further buffered by two further gates in IC1 before going to the drive transistors to ensure that loading from the transistors does not affect anything else. The two buffered outputs of the divider drive the bases of the two drive transistors (Q1, Q2) via resistors R5 and R6, both of which have 470 pF capacitors in parallel with them to ensure fast switching.

The toroidal drive transformer consists of 120 turns around a ferrite core for the primary, and only 10 turns for the secondary. The primary is tapped exactly at the centre to ensure that the carefully derived symmetric drive produces a symmetric current in the core and no saturation problems. The low voltage secondary is used to excite the strain gauges.

OUTPUT AMPLIFIER

The strain gauge output amplifier is a straightforward two-stage amplifier using a common dual operational amplifier with FET inputs. Each amplifier stage is configured as a non-inverting gain stage with the resistor to ground from the negative input capacity coupled so that each gain stage has unity for low frequencies and dc.

The first stage has a gain of about twenty and is adjustable via RV3 to get the signal well above any noise and allow for a span (full-scale reading) adjustment.

The second stage has an identical feedback network to ground but has switchable feedback resistors from the output to the inverting input to set the three weight range gains. For the heaviest weight range the gain of the second stage is set to just a fraction above unity (to be precise, 1.01) by selecting only a 10 ohm resistor in the feedback. This is in parallel with the 100k resistor R15, which is permanently in the feedback. For the middle range (2 kg) a 10k resistor is switched in parallel with R15 to give an overall resistance of 9.09k. This feedback resistance combined with the 1k to ground (R14, via C9) give the stage a gain of 10.09 or near enough to 10.1. Finally, for the most sensitive range, only 100k resistor is left in circuit. Simple number shuffling will show that the gain in this case is 101 times. Hence, exactly 20 dB steps in gain have been given by using only readily available x10 resistors. The scales may only be calibrated on one range so if any gain variation is found it will be necessary to adjust the values of R16 and/or R17 by paralleling or series-ing resistors. If 1%, or better, 0.5%, resistors are used then no problem is expected.

SYNCHRONOUS SWITCH

The output from the ac amplifier and one side of the bridge drive are both identically converted to dc signals by a synchronous switching process using CMOS analogue switches. The CMOS switches used are CD4053BEs which are three separate single-pole double-throw switches implemented in the CMOS process. Separate 47k resistors, R18 and R19, take the two signals to the inputs of two of the switches (the third is not used). All four outputs of the two switches have 1 μF capacitors, C10 to C13 inclusive, in parallel with them to ground to filter the dc that is recovered by the synchronous switching process.

The control lines that drive the analogue switches are derived from the same line that generates the bridge drive voltage. When one half of the bridge is selected the two outputs will be at one extreme of their excursion and when the other half is selected the two outputs will be at the other. Thus, capacitors C10 and C11 will build up a differential voltage exactly proportional to the ac amplifier output and capacitors C12 and C13 have a differential voltage equal to the bridge drive voltage.

The two operational amplifiers in IC6 form two differential amplifiers whose inputs are the two differential signals from the synchronous switching. Capacitors C12 and C13 are fed directly to the differential amplifier but the outputs of C10 and C11 (the strain gauge signal) are buffered first by unity gain non-inverting op-amps to minimise leakage. The gain of the differential amp for the drive signal is only unity as there is plenty of input level available but the gain of the transducer signal is increased to a little over two to optimise noise and stability performance (the exact value of the gain is chosen more by the value of resistors readily available than by anything else).

For optimum common mode rejection of the diff-amps, resistors R20 to R27 should be 0.5% but in practise 1% would be fine. The output of IC6 (pin 1) is a dc voltage exactly proportional to the input drive ac voltage and the output of IC6 (pin 7) is a voltage exactly proportional to the amplified input from the strain gauges, with both outputs having the same constant of proportionality. This is exactly what is required for the DVM to function correctly.

DVM AND DISPLAY

The DVM chip is a National Semiconductor ADD3501/74C935N which uses a pulse width modulation technique to convert the input dc to a digital reading and then generates all the drive and strobe signals necessary to drive a 3½-digit seven segment LED display. The internal clock for all processing is generated by on-chip gates and the external resistor R28 and capacitor C22. The actual conversion is done by switching one end of the resistor R30 between ground and the reference voltage. As the node between R30 and C21 is one input of a comparator and the other input is the input voltage to be converted, and the whole conversion process consists of keeping these two inputs as near to equal as possible by switching the other end of R30 as described. The only way they can be made equal is for

$$\frac{T_{on}}{T_{on}+T_{off}}=\frac{C_{in}}{V_{ref}}$$

The pulse width modulated train thus generated is used to gate the system clock to a system of counters that accumulate the desired reading.

The chip contains all necessary decoders to break the 3½-digits down to seven segment display control lines plus output drivers with sufficient capability to handle the segment lines. Strobe lines out are also provided but these have to be buffered to handle the drive currents required. The four digit-strobe lines are brought out on pins 21 to 24 and are taken through 1k resistors R33 to R36 to the four digit-drive transistors, Q3 to Q6. The collectors of the transistors are taken directly to the common cathodes of the LED displays. The outputs of the DVM are already inverted so they can drive transistors directly. The segment lines have sufficient drive capability to be taken directly to the display and pins 3 to 6 and 26 to 28 carry these signals.

The input analogue signal is filtered by R31 and C20 which are the same value as R30 and C21, the converter mark-space oscillator analogue components. Both are referred to the same analogue ground to minimise offset and noise problems.

Power for the DVM chip is applied to pin 1 for the digital section and via R32 to pin 2 for the analogue section.

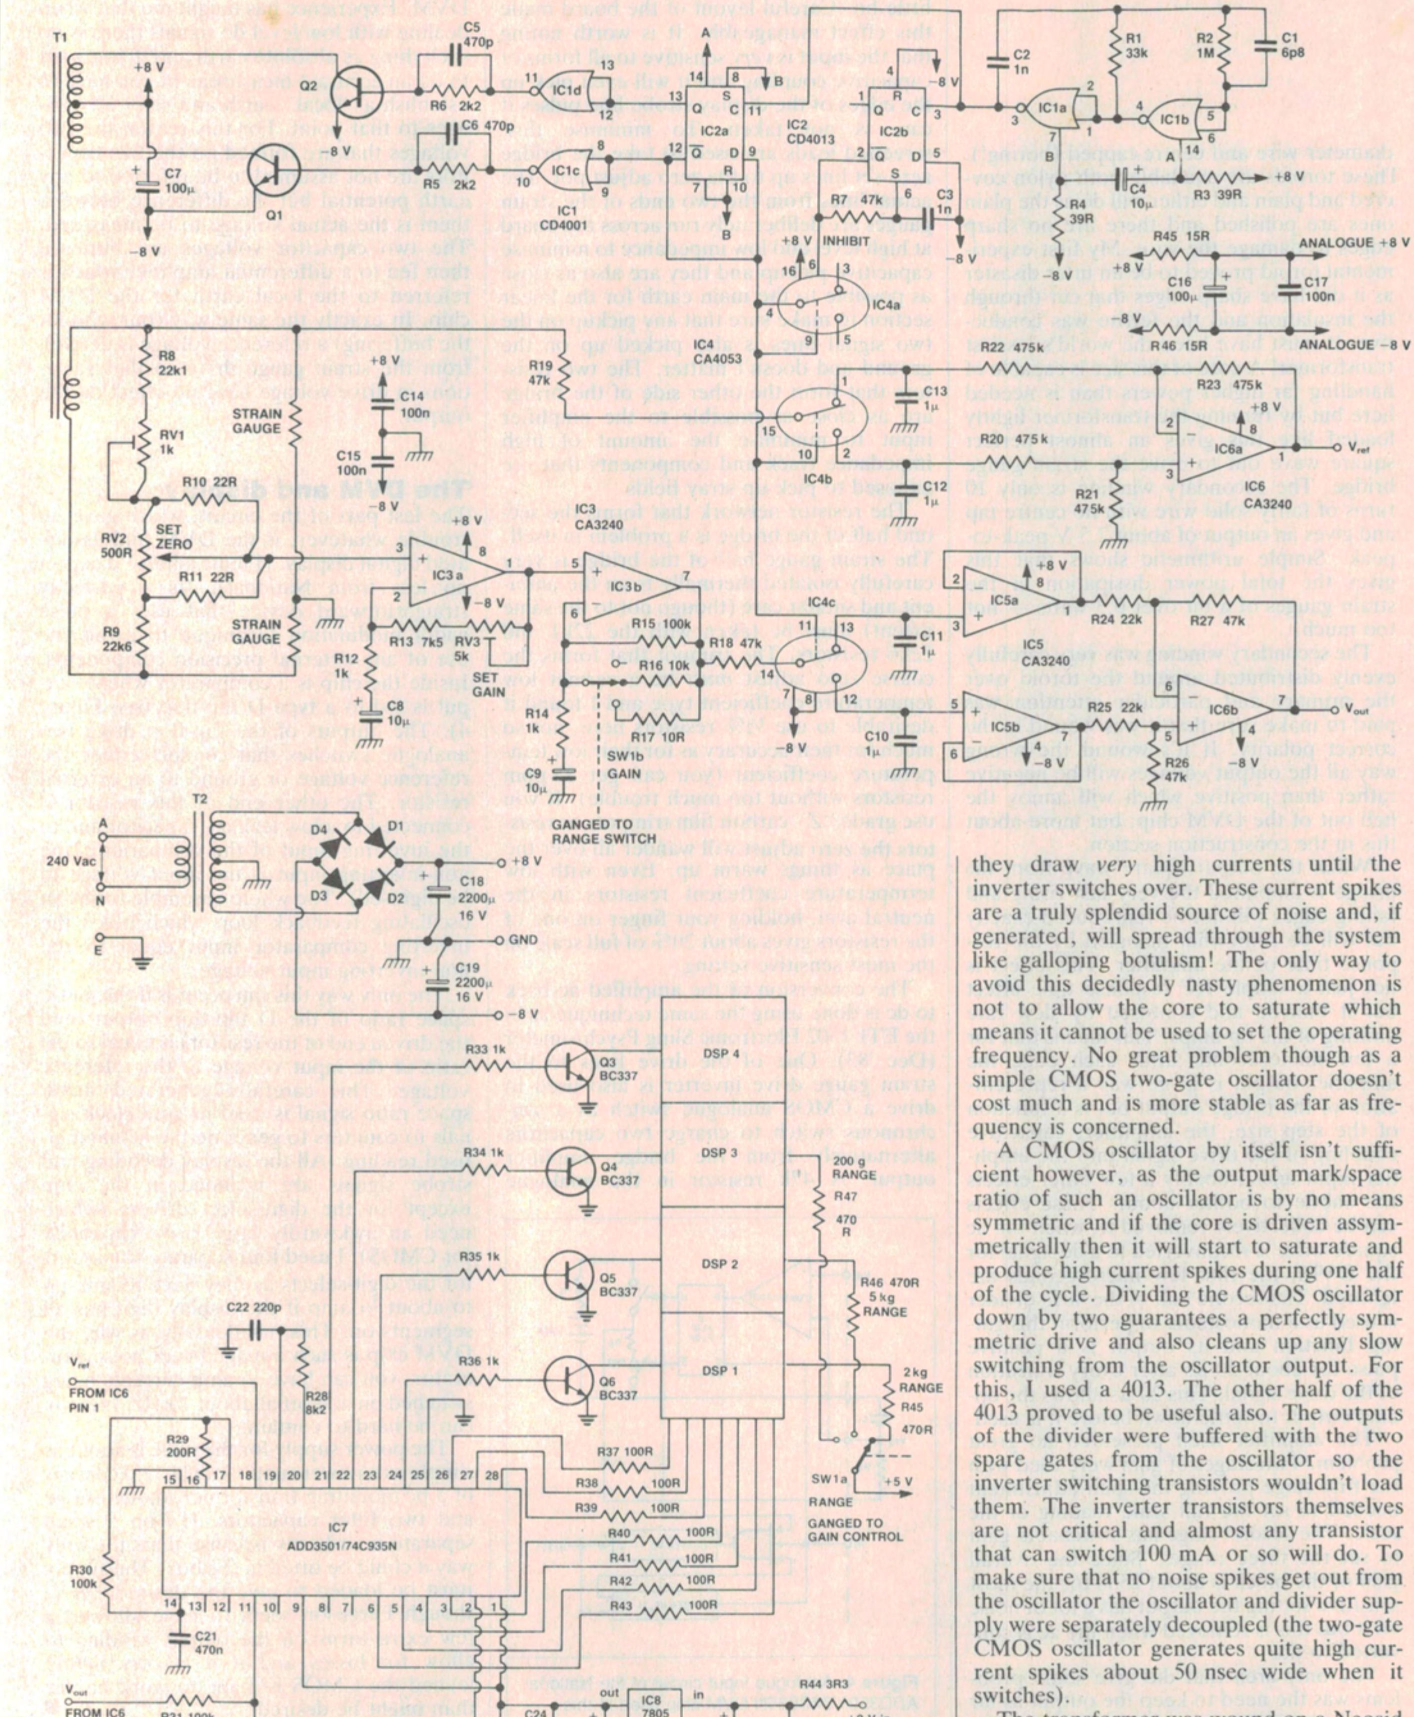

they draw *very* high currents until the inverter switches over. These current spikes are a truly splendid source of noise and, if generated, will spread through the system like galloping botulism! The only way to avoid this decidedly nasty phenomenon is not to allow the core to saturate which means it cannot be used to set the operating frequency. No great problem though as a simple CMOS two-gate oscillator doesn't cost much and is more stable as far as frequency is concerned.

A CMOS oscillator by itself isn't sufficient however, as the output mark/space ratio of such an oscillator is by no means symmetric and if the core is driven assymmetrically then it will start to saturate and produce high current spikes during one half of the cycle. Dividing the CMOS oscillator down by two guarantees a perfectly symmetric drive and also cleans up any slow switching from the oscillator output. For this, I used a 4013. The other half of the 4013 proved to be useful also. The outputs of the divider were buffered with the two spare gates from the oscillator so the inverter switching transistors wouldn't load them. The inverter transistors themselves are not critical and almost any transistor that can switch 100 mA or so will do. To make sure that no noise spikes get out from the oscillator the oscillator and divider supply were separately decoupled (the two-gate CMOS oscillator generates quite high current spikes about 50 nsec wide when it switches).

The transformer was wound on a Neosid #28-680-36 toroid with 120 turns of 14 mil ▶

diameter wire and centre-tapped (boring!). These toroids are available both nylon covered and plain and either will do as the plain ones are polished and there are no sharp edges to damage the wire. My first experimental toroid proved to be an utter disaster as it did have sharp edges that cut through the insulation and the ferrite was conductive. It must have been the world's lossiest transformer! A core of this size is capable of handling far higher powers than is needed here but by running the transformer lightly loaded like this gives an almost perfect square wave out to drive the strain gauge bridge. The secondary winding is only 10 turns of fairly solid wire with no centre tap and gives an output of about 2.5 V peak-to-peak. Simple arithmetic shows that this gives the total power dissipation in the strain gauges of a bit over 0.3 watts — not too much.

The secondary winding was very carefully evenly distributed around the toroid over the primary and particular attention was paid to make sure that it was wound in the correct polarity. If it's wound the wrong way all the output voltages will be negative rather than positive which will annoy the hell out of the DVM chip; but more about this in the construction section.

When the output square wave from the bridge is amplified the very fast rising and falling edges will not be reproduced exactly but will be somewhat mangled by the response time of the amplifier. This effect is not just a simple RC response time effect but it blurred and confused by slew rate limiting of the op-amps. This means that for a few microseconds after each edge the amplifier output is in no way a representation of the bridge output but is a function of the step size, the amplifier, capacitive coupling of the drive signal into the amplifier input and probably a few other effects too subtle to bother about. These effects can be collectively removed if, when the ac square wave is reconverted to a dc signal for the DVM, the first few microseconds are ignored. The second half of the 4013 divider is used as a monostable to perform this gating function and its output is a positive pulse for about 20 µs after every transition of the drive signal. This 20 µs allows the op-amp output to settle down before it is used.

The amplifier itself presented no great problem. Two stages of gain were used with the first stage proving the span adjustment needed to set the full scale reading of the scale. The second stage has switchable gain to set the three ranges. Since the overall gain of the block is about 2500 for the most sensitive setting the output has a lot of noise on it but this is removed when the ac is converted to dc.

The only area that did give some problems was the need to keep the output as far away physically as possible from the extremely sensitive input to the amplifier. It only takes a few *femtofarads* (10^{-15} farads) of capacitive feedback for the whole thing to take off, which doesn't help accuracy one little bit. Careful layout of the board made this effect manageable. It is worth noting that the input is *very* sensitive to all forms of capacitive coupling and it will even pick up the edges of the display strobe line pulses if care is not taken. To minimise this, screened leads are used to take the bridge zero-set lines up to the zero adjust pot. The actual lines from the two ends of the strain gauges are deliberately run across the board at high level and low impedance to minimise capacitive pickup and they are also as close as possible to the main earth for the linear section to make sure that any pickup on the two signal lines is also picked up on the ground and doesn't matter. The two resistors that form the other side of the bridge are as close as possible to the amplifier input to minimise the amount of high impedance track and components that are exposed to pick up stray fields.

The resistor network that forms the second half of the bridge is a problem in itself. The strain gauge half of the bridge is very carefully isolated thermally from the ambient and similar care (though not to the same extent) must be taken with the 22k1 and 22k6 resistors. The trimpot that forms the coarse zero adjust *must* be a cermet low temperature coefficient type and I found it desirable to use ½% resistors here, not so much for their accuracy as for their low temperature coefficient (you can get 25 ppm resistors without too much trouble). If you use grade "Z" carbon film trimpots or resistors the zero adjust will wander all over the place as things warm up. Even with low termperature coefficient resistors in the neutral arm, holding your finger on one of the resistors gives about 20% of full scale on the most sensitive setting.

The conversion of the amplified ac back to dc is done using the same technique as in the ETI-1502 Electronic Sling Psychrometer (Dec '83). One of the drive lines to the strain gauge drive inverter is also used to drive a CMOS analogue switch as a synchronous switch to charge two capacitors alternatively from the bridge amplifier output. A 47k resistor in the analogue switch, combined with the two capacitors ters out the noise from the op-amps and gives a clean and stable dc signal for the DVM. Experience has taught me that when dealing with low level dc signals there is no such thing as absolute earth and if you want to do an accurate measurement you have to establish a "local" earth and refer all voltages to that point. For this reason the two voltages that are formed on the two capacitors are not assumed to be referred to any earth potential but the difference between them is the actual voltage to be measured. The two capacitor voltages are buffered then fed to a differential amplifier which is referred to the local earth for the DVM chip. In exactly the same way (but without the buffering) a reference voltage is derived from the strain gauge drive so that variations in drive voltage have no effect on the output.

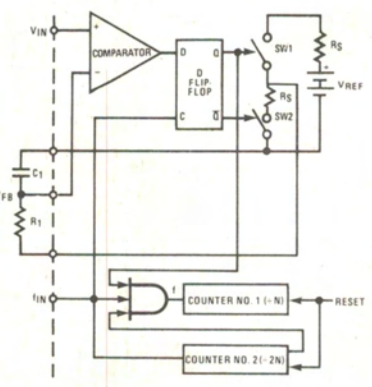

Figure 4. Analogue input circuit of the National ADD3501/74C935N DVM chip used in this project.

The DVM and display

The last part of the circuit, which gave no trouble whatever, is the DVM and associated digital display. (I must admit I'd expect no less from National). Its a perfectly straightforward device that uses a pulse width modulation technique to avoid the use of any external precision components. Inside the chip is a comparator whose output is fed to a type-D flip-flop (see Figure 4). The outputs of the flip-flop drive two analogue switches that connect either the reference voltage or ground to an external resistor. The other end of this resistor is connected to a low leakage capacitor and to the inverting input of the comparator (the non-inverting input is the input voltage to be digitised). The whole ensemble forms an oscillating feedback loop which holds the inverting comparator input equal to the non-inverting input voltage.

The only way this can occur is if the mark/space ratio of the D flip-flop output (and the driven end of the resistor) is equal to the ratio of the input voltage to the reference voltage. This carefully generated mark/space ratio signal is used to gate clock signals to counters to generate the actual digitised reading. All the display decoding and strobe signals are included in the chip except for the digit-select drivers (which need an awkwardly large drive capability for CMOS). I used four separate transistors for the digit-selects as they have to sink up to about ¼ amp if the display digit has all segments on. This, incidentally, is why the DVM chip is such a magnificent noise generator; you can have ¼ amp currents being switched on and off at about 1 kHz. It really can be hard to contain.

The power supply for the scale is about as simple as a power supply can be and consists of a pc mounting transformer, diode bridge and two filter capacitors. It's on a small separate board only because thats the only way it could be fitted in the box. The supply must be loaded to give the desired ±8 V, though Ferguson seem to have allowed a few extra turns on the output winding to allow for losses and if it is very lightly loaded the CMOS is made to work harder than might be desired. ●

. . . continued next month