Brim Full

Capacitive liquid-level measurement

Wolfgang Rudolph (Germany), Rudolf Pretzenbacher (Austria), and Burkhard Kainka (Germany)

Electronics enthusiasts are sometimes a breed apart. Most people simply look at a bottle when they want to know how full it is, but we want to measure it.

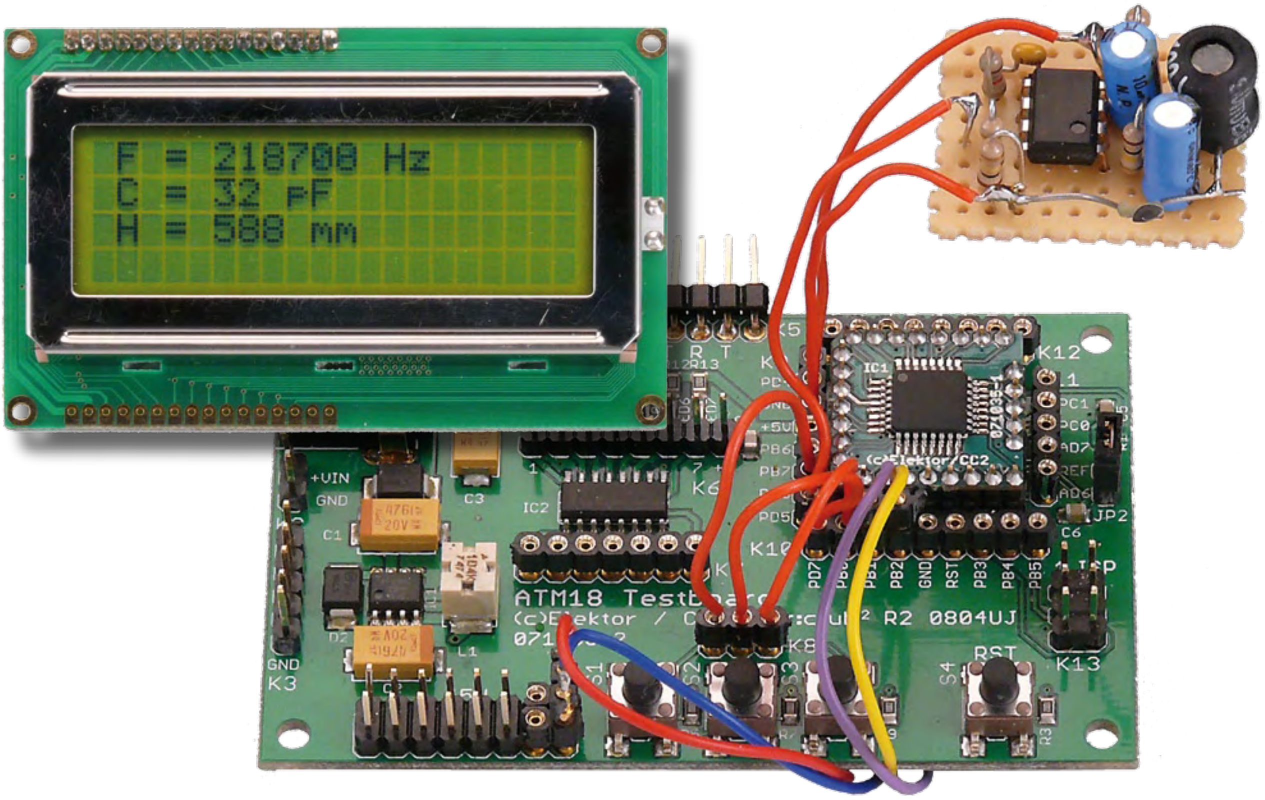
```
F = 218708 Hz
C = 32 pF
H = 588 mm
```

Of course, it doesn't have to be a bottle. Situations that involve measuring the level of a liquid stir the creative juices and foster true acts of genius, and there are countless applications for liquid-level sensors, ranging from rain barrels to heating-oil tanks.

We're sure that our readers can come up with many other situations where the liquid-level sensor described here can be put to good use. However, let's first consider the question of how to measure a liquid level accurately and reliably.

Measuring methods

A wide variety of measuring methods are used. Many lavatory cisterns have a float valve that first reduces the inflow of water when the float rises to a certain level and finally stops it completely. In

Table 1

Inductor specifications

(vertical package
with moderate rated current)

Manufacturer: Fastron;
type number 09 P-103 J-50

Dimensions: Ø 9.5 mm, height 14 mm,
lead pitch 5 mm

Inductance: 10.0 mH (at 20 kHz)

Self-resonant frequency (SRF): 0.41 MHz

Rated DC current: 90 mA

Resistance: 35.0 Ω

Tolerance: ±5 %

Q (min): 70

this case, the float is not only the sensor but also the actuator, which controls the valve via a lever mechanism. Although this is a very reliable principle, it can't be used to measure the liquid level. The same principle was used in the past (and is sometimes still used) to measure the fuel level in petrol tanks of cars. In this case, the float moves the wiper of a potentiometer instead of actuating a valve. This variable resistance forms part of a voltage divider that drives a milliammeter, which indicates how full the tank is. In some cases, the accuracy of this gauge leaves a lot to be desired.

Nowadays a wide variety of modern measuring methods are used in many different situations. They include hydrostatic and differential pressure measurement, conductivity measurement, light absorption measurement, transit time measurement using ultrasound, distance measurement using microwaves, and even transit time measurement using radar pulses.

From an electronic perspective, capacitive measurement is also interesting. This method involves measuring the change in the capacitance between two electrodes. If these electrodes are located in a container with a liquid that covers them more or less depending on its level, the capacitance of this 'capacitor' changes accordingly. The capacitance depends on the dielectric constant $\varepsilon$of the liquid, and it increases as the level of the liquid rises.

Capacitive sensing

You've probably guessed that this is the method we intend to use here. After all, we're used to working with capacitors. However, it's not as simple as it seems at first glance. We have to do a bit of maths first. This article is based on a capacitive liquid-level sensor built by Rudolf Pretzenbacher, which uses a simple but remarkably stable oscillator for the sensor circuit and an AVR microcontroller for the signal processing. His liquid-level gauge provided the inspiration for this ATM18 article, and it delivers truly astounding results. This setup can be used to measure capacitances in the range of nanofarads (nF) to femtofarads (fF). In case you've forgotten, a femtofarad is 10^{-15} F or a thousandth of a picofarad. How can such high sensitivity be achieved? The answer is that the 'sense capacitor' in the liquid is one of the frequency-determining components of a resonant loop, which in turn is part of an oscillator circuit. If an object to be measured is brought in the vicinity of the capacitor, the resonant frequency of the loop changes. The more the capacitance of the capacitor is increased by the object, the lower the resulting frequency. The task of the microcontroller on the Elektor ATM18 board is to measure the frequency and then calculate the value of the capacitance from the measured frequency and the known value of the inductance.

This sounds quite simple, but there are still a few details to be sorted out.

Oscillator

The oscillator circuit can affect the resonant loop due to its own capacitance or as a result of excessively strong coupling. To keep this effect as small as possible, the resonant loop should have a high quality factor (Q) and the excitation level should be kept low. It is also important to choose a suitable inductor.

In this case, we decided on a fixed inductor made by Fastron. This inductor (type number 09 P-103 J-50; available from Reichelt and other sources) has an inductance of 10 mH, a DC resistance of 35 Ω, and a self-resonant frequency of 410 kHz. This means that it has a remarkably low stray capacitance of 15 pF. In addition, it has a specified Q factor of 70 (max.). Its characteristics are listed in **Table 1.**

The higher the Q factor of a resonant loop, the lower its damping. A Q factor of 70 means that the amplitude of a 'free' (damped) oscillation is reduced by a factor of e after 70 cycles, which can be seen very nicely on an oscilloscope. The damping results from the resistive losses in the wire and the magnetic losses in the core. A resonant loop with an inductance of 10 mH and a capacitance of 6300 pF has a resonant frequency of 20 kHz, and the inductive and capacitive impedance are both 1260 Ω. The ratio of this impedance to the DC resistance (35 Ω) yields a theoretical Q factor of 36, which means that the resonant impedance of the circuit is 45 kΩ (1260 Ω × 36). The Q factor and the resonant impedance increase as the capacitance is reduced and the frequency rises. For a high Q factor, we have to aim for a high L/C ratio. At around 3000 pF and 30 kHz, the calculated value of the Q factor is approximately 70. The core losses increase at very high frequencies, which causes the Q factor to drop. However, the oscillator circuit has an even larger effect, since a resonant loop with a high resonant impedance is especially sensitive to external influences.

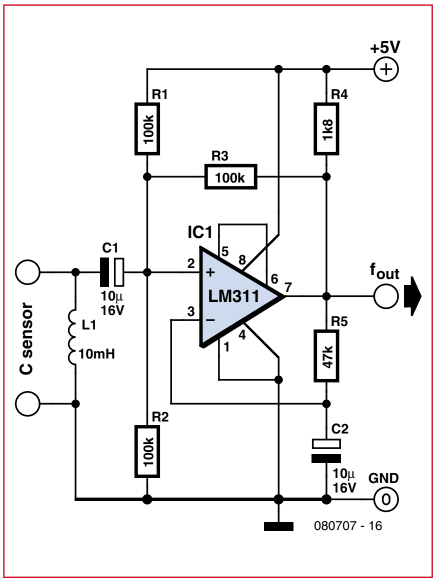

Figure 1. Schematic diagram of the oscillator used for capacitance measurement.

Figure 1 shows the oscillator circuit used here, which is built around an LM311 comparator. It compares the input voltage with a reference voltage and converts the sinusoidal signal from the resonant loop into a squarewave signal at its output. This signal excites the resonant loop via a feedback resistor. A voltage divider at the non-inverting input of the comparator provides a voltage equal to half the supply voltage. The inverting input is fed by a comparison voltage obtained by integrating the output voltage. As a

Listing 1

Capacitance measurement

```
Config Timer0 = Timer ,
   Prescale = 64
Config Timer1 = Counter ,
   Edge = Falling , Prescale
   = 1
On Ovf0 Tim0_isr
On Ovf1 Tim1_isr
Enable Timer0
Enable Timer1

Do
  Ticks = 0
  Enable Interrupts
  Waitms 1100
  Disable Interrupts
  Lcdpos = 2 : Lcdline = 1 :
   Lcd_pos
  Lcdtext = "Freq = "
  Lcdtext = Lcdtext +
   Str(freq)
  Lcdtext = Lcdtext + " Hz
   "
  Lcd_text
  Print Freq;
  Print " Hz"
  C = Freq / 10000000
  C = 1 / C
  C = C * C
  C = C / 39.48
  If Pinb.0 = 0 Then C0 = C
  C = C - C0
  Print Fusing(c , "#.###");
  Print " pF"
  Lcdpos = 2 : Lcdline = 2 :
   Lcd_pos
  Lcdtext = "Cap ="
  Lcdtext = Lcdtext +
   Fusing(c , "#.###")
  Text = Fusing(c , "#.###")
  Lcdtext = Text
  Lcdtext = Lcdtext + " pF
   "
  Lcd_text
  Waitms 10
Loop


Tim0_isr:
  '1000 µs
  Timer0 = 6
  Ticks = Ticks + 1
  If Ticks = 1 Then
    Timer1 = 0
    Highword = 0
  End If
  If Ticks = 1001 Then
    Lowword = Timer1
    Freq = Highword * 65536
    Freq = Freq + Lowword
    Ticks = 0
  End If
Return

Tim1_isr:
  Highword = Highword + 1
Return
```

result, the operating point of the oscillator is set automatically, and it starts reliably and produces a symmetric square wave at the output.

With regard to the effect of the oscillator circuit on the resonant loop, the main consideration is the resistor values. The voltage divider formed by the two 100-kΩ resistors loads and thus damps the resonant loop with an effective value of 50 kΩ. There is also the resistance of the negative feedback resistor (100 kΩ) divided by the effective voltage gain. As a result, stable oscillation is possible with sensor capacitance values of up to 100,000 pF (or more). The open-circuit frequency is approximately 350 kHz, which yields an effective capacitance of around 20 pF. The inductor accounts for 15 pF of this, while the input capacitance of the LM311 and the stray circuit capacitance add another 5 pF.

If you use an oscilloscope to view the signal on the inductor, you will see an amplitude of approximately 1 V at the highest frequency and a somewhat distorted sinusoidal waveform. This means that the excitation level could be reduced even further. However, with increasing sensor capacitance the amplitude decreases noticeably and the signal becomes more sinusoidal. The oscillator still works at 100 nF, with a frequency of 4.9 kHz and a signal amplitude of 0.1 V. It stops operating suddenly somewhere above this figure.

The next issue to be considered is frequency stability. The fact that the circuit only contributes 5 pF to the capacitance of the resonant loop is in itself favourable. This leaves us with the difficult question of the temperature dependence of the inductance. The only way to answer this question is to perform experiments. To make a long story short, we can say that the stability of the prototype version built on stripboard in the Elektor labs (**Figure 2**) is sufficient to achieve a sensitivity of 0.001 pF, or in other words 1 fF (1 femtofarad – what an uncommon term!). Incidentally, frequency measurement is not the limiting factor. At 350 kHz and 20 pF, a change of 1 Hz corresponds to a capacitance change of only around 0.1 fF. However, the effective constancy is somewhat lower.

Frequency measurement

Now we come to familiar ground. Frequency measurement was already described in instalment 4 of the Bascom AVR series (Elektor December 2008). The counter input is T1 (PD5), and the frequency in hertz can be obtained directly with a gate period of 1 second. It is sent directly to the PC at 9600 baud, without any correction or window dressing. All that's left is to convert the frequency into capacitance. We use a single-precision variable for this. The conversion formula must be broken down into individual operations in Bascom. Here you have to ensure that the intermediate values do not become too large or too small, since this would degrade the accuracy. This means that the sequence of the operations is somewhat important. The 10 mH of the inductor is expressed as a factor of 10,000,000. The underlying reason for this is to arrive at a value in picofarads at the end. If comparative measurements indicate that the actual value of the inductor is slightly different, such as 1% higher or lower, this is the place to make the correction. The inductor has a rated tolerance of 5%, which means that the capacitance can be measured with a potential error of approximately 5%.

Body capacitance

If you move your hand close to the oscillator (Figures 1 and 2), you will see the measured capacitance change by a few femtofarads, even if no sensor cable is connected. We measured the following approximate results at various distances between the board and our hand:

5 cm	0.005 pF
4 cm	0.009 pF
3 cm	0.020 pF
2 cm	0.040 pF
1 cm	0.100 pF

This is interesting from a physics perspective. The phenomenon of body capacitance is both familiar and notorious among radio hobbyists. If a DIY receiver is not adequately screened, it is often possible to detune it slightly by moving your hand toward it. Some people make handy use of this effect for fine tuning when receiving SSB signals.

Musicians who use Theremin instruments also take advantage of body capacitance.

The open-circuit capacitance C_0 is around 20 pF. Of course, the exact value depends on several factors, including component tolerances, PCB construction, and perhaps even the type of solder that is used, since the dielectric constant of solder flux can have an effect on the order of a few femtofarads. The only solution to this is to perform a zero-point calibration.

Nothing could be easier: when the user presses a button connected to port B0, the current zero-point capacitance C_0 is measured and stored. This is anyhow necessary, because if you use a cable to connect the sensor it can easily contribute another 10 pF. Consequently, we measure and store the zero offset before making the actual measurement, and this way we obtain the best possible accuracy

The measured values are output in two different ways: via the serial interface and on the familiar LCD with its two-wire interface. At first this was a bit too much for the LCD routine, which didn't want to cooperate with the timer interrupts. The problem was found to arise from passing variables to the subroutines, and it was cured by declaring all variable as global. In addition, the timing was improved to make data transfer even more reliable (see **Listing 1**).

Their hand movements alter the frequency of an oscillator and thus change the audio frequency in a smooth, continuous manner.

You can try this for yourself with this oscillator. Connect a copper-plated board in Eurocard format (100× 160 mm) to act as the sense electrode. This adds approximately 17 pF to the capacitance of the resonant loop, and the frequency drops to around 260 kHz. This is in the long-wave radio band, and you can pick up the signal on a radio. With a bit of luck, you can find a long-wave broadcast signal that interferes with the oscillator signal to produce a beat frequency. Then you can start making music, assuming you have the knack.

All the neighbourhood cats will probably run for cover, but that shouldn't stop you from trying out the effect and learning to understand it, even if you'll never compete with Theremin virtuoso Lydia Kavina, a great-niece of the inventor of the Theremin. The most effective variation in capacitance, around 0.1 pF, occurs at a distance of around 5 cm due to the relatively large size of the sense electrode.

Now the program displays the current frequency and the capacitance. This enables us to make some experimental measurements of temperature stability. For example, you can warm the inductor with your hand and observe the change. With a temperature increase of approximately 20 °C (to around 30 °C), the measured capacitance increased by approximately 0.15 pF. This means that if your objective is to measure the value of an unknown capacitor, the temperature is scarcely important. However, if you actually want to measure capacitance with an accuracy of a few femtofarads, you must first allow the oscillator to stabilise for a few minutes and then make a zero-offset reading. The measured value changes by less than 5 fF over the course of several minutes.

Capacitance measurement

People who play around with RF circuits almost always have something to measure, such as a variable capacitor. Before a true radio hobbyist tosses an old radio in the bin, he at least salvages the variable capacitor, since they are not so easy to come by nowadays. Naturally, you have to measure the salvaged part to know what you actually have. If it has a range of 8 pF to 520 pF, it's brilliant.

You can also measure unknown SMD capacitors, variable-capacitance diodes, the input capacitances of FETs or valves, and cable capacitances. You can even determine the length of a cable by measuring its capacitance. For example, suppose you have a partially used roll of coax cable and you want to feed it down a disused chimney. Before you start, it's a good idea to know whether it's long enough to reach the bottom. We've all heard enough stories about cursing men on high roofs.

This question is easily answered with our capacitance meter. The capacitance per metre is stated on the data sheet. For example, popular 50-Ω RG58 cable has a capacitance of 100 pF/m. If you don't have a data sheet, you can simply measure the capacitance of a known length, such as 1 metre, to determine the number of picofarads per metre. Once you know this value, you can easily calculate the cable length from the measured cable capacitance (cable capacitance divided by capacitance per metre yields cable length in metres). The fact that the cable also has an inductance doesn't matter, since the measuring frequency is much less than the quarter-wavelength frequency. For example, at 100 kHz the wavelength is 3 km.

Liquid level measurement

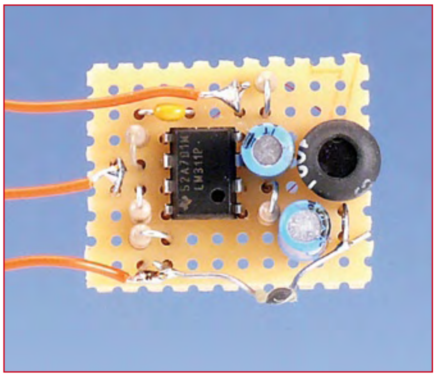

Figure 2. Prototype version of the oscillator, built on a piece of perforated circuit board.

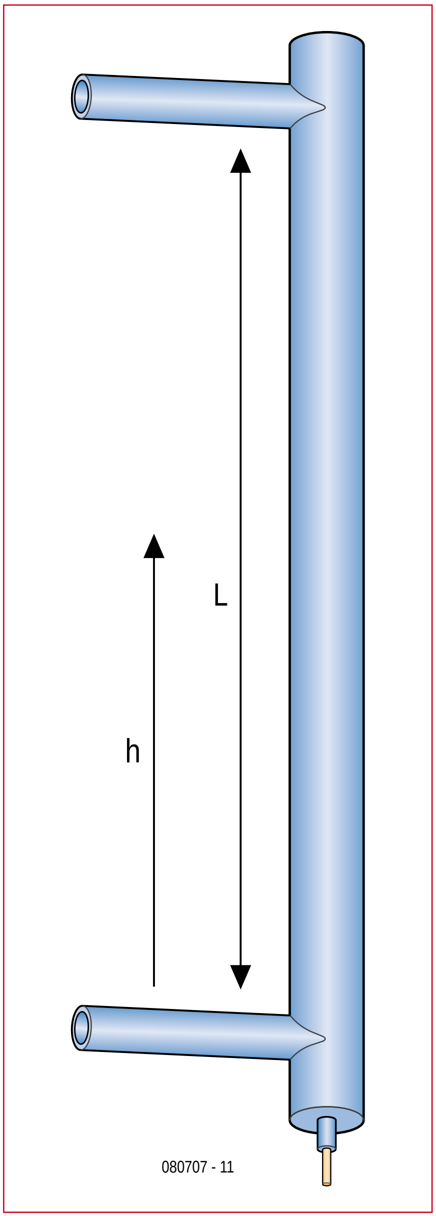

Figure 3. The liquid-level sensor is a tube with an insulated inner electrode that forms a cylindrical capacitor. Here L is the length of the active portion of the tube (wrapped with aluminium foil) and h is the height of the water in the tube.

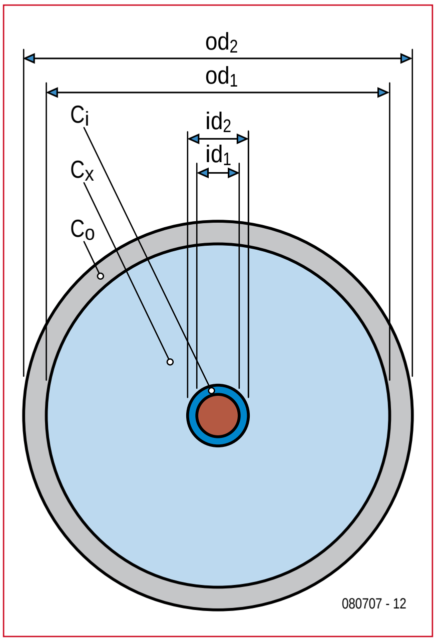

Figure 4. The concentric capacitors of the sensor tube structure.

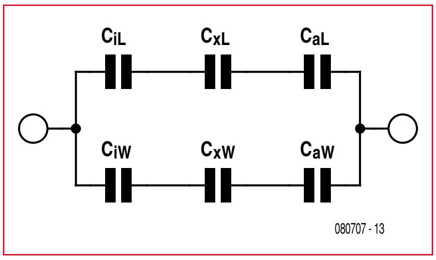

Figure 5. The equivalent circuit of the sensor tube.

To make our liquid-level sensor, we fitted a small Plexiglas (polycarbonate) tube with two connection stubs. A length of polyethylene-insulated hookup wire was stretched through the tube and centred as well as possible, and then both ends of the tube were sealed watertight (**Figure 3**). The conductor of the hookup wire must be fully insulated (galvanically isolated) from the space inside the tube. Then we wrapped the length of the tube between the two stubs with aluminium foil applied as uniformly as possible and attached a bare connecting lead to the aluminium foil (held in place by electrician's tape). The bare lead and the end of the hookup wire protruding from the tube form the terminals of our sense capacitor.

A cylindrical capacitor is a rotationally symmetric form, so its capacitance can be calculated rather accurately by using the following formula if the length is much greater than the diameter:

$$c = \frac{2 \cdot \pi \cdot \varepsilon_0 \cdot \varepsilon_r \cdot l}{LN\left(\frac{od}{id}\right)}$$

ε_0 = dielectric constant of vacuum and air (8.854×10^{-12} As/Vm)
ε_r = relative dielectric constant (material constant)
L = cylinder length
od = diameter of the outer electrode (here od2)
id = diameter of the inner electrode (here id1)

If we combine the constants and convert metres to millimetres, we obtain the following formula:

$$c = \frac{0.0556 \cdot \varepsilon_r}{LN\left(\frac{od}{id}\right)} \cdot l \quad \text{pF/mm}$$

If a cylindrical capacitor consists of several concentric layers, each layer forms a separate capacitor (here C_o, C_x, and C_i). The total capacitance is then determined by the series connection of the individual capacitors (**Figure 4**). If we divide the cylindrical capacitor into a portion filled with water or another liquid (C_W) and a portion filled with air (C_A), the total capacitance of the tube is $C_T = C_W + C_A$ (parallel connection), with the portion filled with water having a length h and the portion filled with air having a length L – h. The equivalent circuit of this arrangement is shown in **Figure 5**.

The relative dielectric constant (ε_r) of air is 1.0, while the relative dielectric constant of water depends on the temperature and ranges from 55 to 88 (approximately 83 at 10 °C). The dielectric constant of transparent plastic is around 3.0 (polystyrene and polycarbonate) or 3.2 (acrylic), and the dielectric constant of wire insulation is around 2.3 (polyethylene) or 4 to 5 (polyvinyl chloride).

This is excellent for our intended measuring applications because it means that there will be a rather large difference between the values of the capacitance Cx in air and in water.

The capacitances in the air-filled portion of the tube are:

$$CiA = \frac{0.0556 \cdot 2.3 \cdot (l-h)}{LN\left(\frac{id2}{id1}\right)}$$

$$Cxl = \frac{0.0556 \cdot 1 \cdot (l-h)}{LN\left(\frac{od1}{id2}\right)}$$

$$CoA = \frac{0.0556 \cdot 3 \cdot (l-h)}{LN\left(\frac{od2}{od1}\right)}$$

Table 2

Sensor tube data (for Figure 6)

Standpipe outside diameter: 12 mm

Standpipe inside diameter: 8.5 mm

Standpipe length: 300 mm

Inner electrode
conductor diameter: 0.4 mm

Inner electrode
outside diameter: 0.6 mm

Standpipe tube dielectric constant: 3.0

Inner electrode dielectric constant: 2.3

Electrolyte dielectric constant: 83

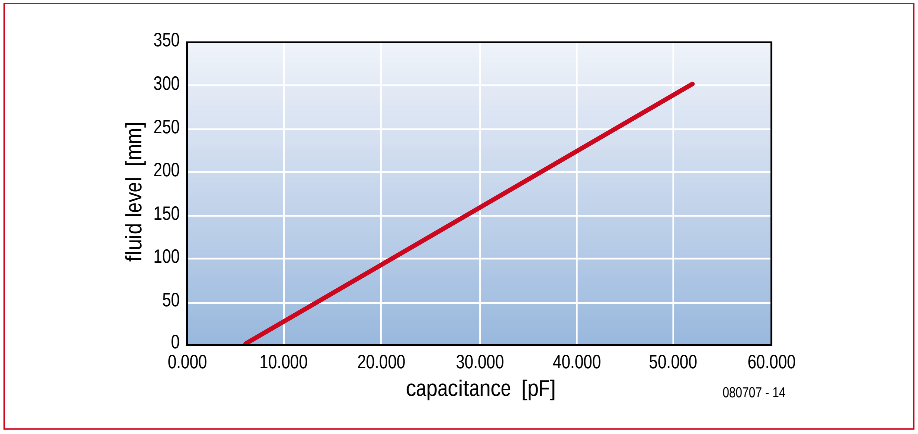

Figure 6. The capacitance increases linearly with the liquid level.

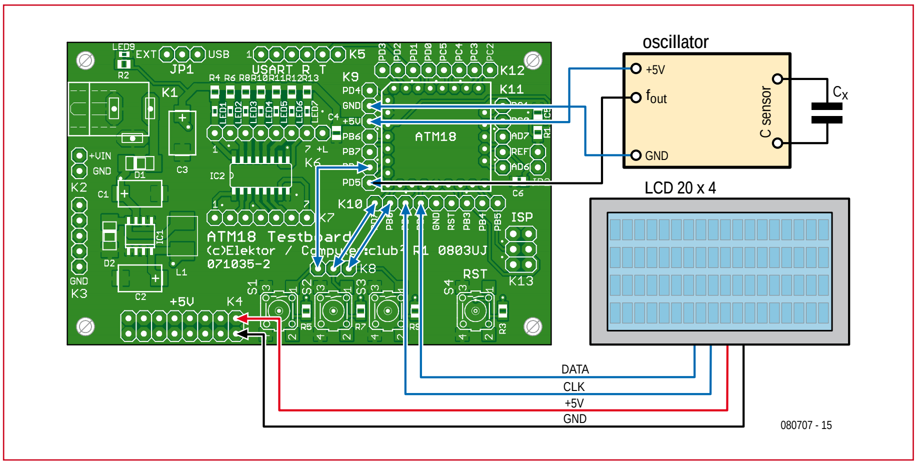

Figure 7. Wiring diagram of the Elektor ATM18 board for the liquid-level gauge.

while the capacitances in the water-filled portion are:

$$CiW = \frac{0.0556 \cdot 2.3 \cdot h}{LN\left(\frac{id2}{id1}\right)}$$

$$CxW = \frac{0.0556 \cdot 83 \cdot h}{LN\left(\frac{od1}{id2}\right)}$$

$$CoW = \frac{0.0556 \cdot 3 \cdot h}{LN\left(\frac{od2}{od1}\right)}$$

If you use a spreadsheet program to calculate and plot the relationship between the total capacitance and the water level, you will discover that it is fully linear if you use a fixed dielectric constant for water. **Figure 6** shows the capacitance as a function of liquid level for a standpipe sensor with the dimensions given in **Table 2**.

Now we can use our standpipe sense capacitor and an inductor with a more or less known value to form a resonant loop, measure the resonant frequency, and use the well-known resonant-loop formula

$$f0 = \frac{1}{2 \cdot \pi \cdot \sqrt{(L \cdot C)}}$$

to calculate the capacitance of the standpipe and thus determine the height of the water in the standpipe. We first measure the capacitance Cmin with the standpipe empty (h = 0) and the maximum capacitance Cmax with the standpipe full (h = L), after which we can use the straight-line formula to calculate the height:

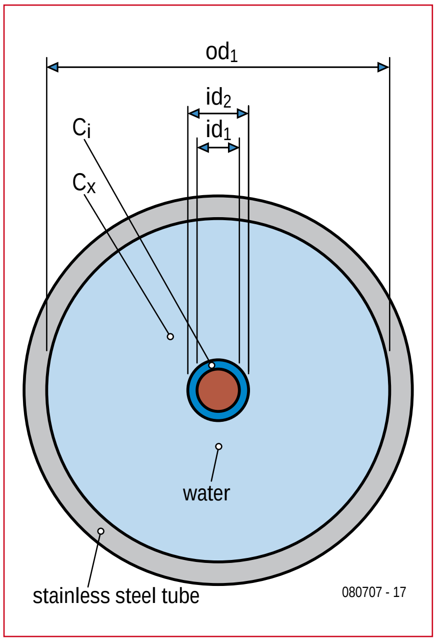

Figure 8. Simplified sensor construction using a stainless-steel or copper outer tube and an insulated brass tube as the inner electrode.

$$h = \frac{L \cdot (C_{measured} - C_{\min})}{C_{\max} - C_{\min}}$$

Here the mechanical accuracy of the construction and the accuracy of the reference inductor do not matter, and the absolute accuracy of the frequency measurement, the presence of parasitic capacitances, and the dielectric constants of the materials used to construct the sensor are equally irrelevant.

The oscillator module (**Figure 2**) should be located as close to the sensor as possible in order to minimise the parasitic capacitance of the cable and reduce the effects of nearby objects on the sensor cable capacitance.

Listing 2

Calibration and calculation of the liquid level

```
Hmin = 0.0
 Hmax = 300.0
 Getminmax
 If Cmax <= Cmin Then
   Cmin = 7.0
   Cmax = 52.0
 End If
…

Sub Calclevel
 ‘ensure that: Hmax>Hmin and
   Cmax>Cmin
 If Cap < Cmin Then Cap = Cmin
 K = Hmax - Hmin
 D = Cmax - Cmin
 If D = 0 Then D = 0.01  ‘avoid
   division by zero
 K = K / D
 D = -k
 D = D * Cmin
 Y = Cap * K
 Y = Y + D
 Yfix = Y
End Sub

‘Calibrate Minimum Value
Sub Calibmin
 Print “Minimum Calibration”
 Bitwait Pind.7 , Set
 Cmin = Cap
 Print “Cmin” ; Cfix ; “ pF”
 Eadr = Eadrcmin
 Writeeeprom Cmin , Eadr
End Sub

‘calibrate Maximaum Value
Sub Calibmax
 Print “Maximum Calibration”
 Bitwait Pind.6 , Set
 Cmax = Cap
 Print “Cmax” ; Cfix ; “ pF”
 Eadr = Eadrcmax
 Writeeeprom Cmax , Eadr
End Sub
```

Software

The Bascom project *Level.bas* also uses the serial interface and the LCD. In addition to the frequency and the capacitance, it shows the liquid level in millimetres on the display. A pair of buttons connected to PD6 and PD7 can be used for calibration, with the

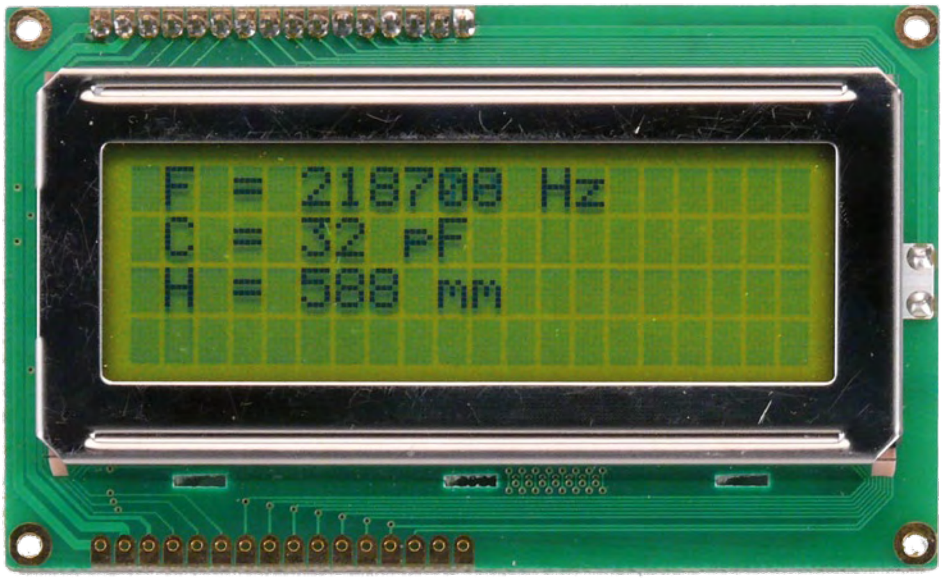

calibration values being stored in EEPROM. The default values assign a height of 0 to a capacitance of 7 pF and a height of 300 mm to a capacitance of 52 pF. If you adjust the liquid level to a height of 0 mm and press the first button (PD7), the measured capacitance is copied to Cmin and stored in memory. After this, you can fill the sensor tube to the 300-mm level and press the second button (PD7) to copy the corresponding value to Cmax. This data is held in non-volatile memory, so it is available the next time you switch on the instrument (see **Listing 2**).

If the parasitic capacitance of the cable (approximately 33 pF) is taken into account, the measured values are amazingly close to the theoretically determined values. From this we can conclude that a method based on purely theoretical calculation (without calibration of the minimum and maximum levels), and taking the temperature dependencies of the electrolytes into account, could be implemented with a reasonable amount of effort.

As already mentioned, the simple approach only works if you assume that the dielectric constant of the electrolyte (in this case water) remains more or less the same after calibration. The error due to electrolyte temperature variation depends on the dimensions of the sensor tube, and with the prototype arrangement it is approximately 1 mm per 20 °C.

If this is not acceptable, you will have to measure the temperature of the electrolyte as well and use a table to determine the actual dielectric constant. Unfortunately, the simple calibration procedure is no longer feasible in this case, and the liquid level must be determined using the theoretical formulae. With this approach, the accuracy of the sensor tube construction, the exactness of the dielectric constants of the tube insulation and the insulation of the centre electrode, and the accuracy of the reference inductor and the frequency measurement are very important for obtaining good results. In addition, the parasitic capacitance of the connecting cable must be measured exactly.

Choice of materials

A wire with polyethylene (PE) insulation is a better choice for the inner conductor than one insulated with polyvinyl chloride (PVC) because the dielectric constant of polyethylene has a very small range of variation and lies between 2.28 and 2.3. A good way to obtain such a wire is to remove the sheath and braid from a length of coax cable. If the dielectric is transparent, it is solid polyethylene with $\varepsilon_r = 2.3$. Naturally, you can also use a glass tube (ε_r range: 6 to 8) for the sensor.

It's even easier if you can allow the electrolyte to make electrical contact with a sensor electrode and the electrolyte is electrically conductive (which is the case with normal water). In this case the electrolyte acts as the outer electrode of the capacitor (see **Figure 8**). Here again there is a linear relationship between the capacitance and the liquid level. The temperature dependence of the electrolyte is largely irrelevant as long as the conductivity of the electrolyte is much greater than the conductivity of the insulation of the inner electrode. This is always the case with tap water.

Constructing the sensor is a bit tricky in this case because the inner electrode cannot be clamped at both ends. The best approach is to use a thin brass tube (from a DIY shop) and insulate it with heat-shrink tubing so the brass does not come in contact with the electrolyte. Now the trick is to devise brackets that hold the inner tube and the outer tube of the sensor (the outer tube can be made from stainless steel or copper) such that they are accurately concentric. Depending on the diameter of the outer tube, an arrangement using plastic champagne corks with a hole drilled through the centre is reasonably effective. Don't forget to also drill a vent hole.

(080707-I)