Experiments with a novel no-moving-parts anemometer

Tim Williams
July 2025

Introduction

I've always had a fascination with recording the weather, not, I think, a very unusual preoccupation, especially in the UK. I've had a Davis Vantage Pro II system on the roof for getting on for 15 years and, with occasional maintenance attention, it's given good service, but now looks rather careworn. At the same time I've been incubating an idea for an alternative method for measuring the windspeed and direction, other than the conventional spinning cups and windvane, and have finally got round to trying to make it work. This article describes the idea, its implementation, and some results.

Of course, the first thing to do when trying out a new idea is to see if anyone else has done it first. I didn't do this, which in hindsight was a good thing; if I had, I'd have found the excellent YouTube video on the Smart Solutions for Home (SSH) channel, which is so beautifully and professionally done that it puts anything I could do completely in the shade. But then, I don't have a laser cutter and 3-D printer at my disposal. That design uses a similar principle but implements it in a different way. My mechanical design is considerably more knife-and-fork but can be done with a minimum of modern workshop equipment if you don't want a precision product.

Having said which, the results of the experimental designs are, to put it politely, somewhat equivocal. I wouldn't recommend doing it exactly this way again. But it seems that the experience gained is worth writing up so that, if anyone else has a similar idea, they don't have to start from ground zero to implement it.

No-moving-parts anemometers are not particularly new, but they are very largely dominated by the ultrasonic type, exemplified by Gill Instruments' Windsonic range. Here the wind is deduced from time-of-flight measurements of an ultrasonic pulse over two orthogonal paths. All that is needed of the sensor is a pair of ultrasonic transducers and some nifty signal processing; nothing moves in the sensor head. A comparison of conventional cup-and-vane anemometers with the ultrasonic variety is given here (provided by an ultrasonic sales company):

Factor	Traditional Anemometers	Ultrasonic Anemometers
Cost	Low upfront cost, high long-term maintenance	High upfront cost, minimal maintenance
Accuracy	Moderate (degrades with wear)	High (stable over time)
Response Time	Slow (mechanical lag)	Instantaneous (real-time updates)
Durability	Prone to ice/dust damage	Rugged; handles extreme temps, storms, and dirt
Installation	Simple, but requires sturdy mounting	Flexible, but needs precise alignment
Environmental Fit	Mild climates, low wind variability	Harsh climates, high-precision needs
Maintenance	Frequent cleaning, lubrication, part replacement	Rare calibration, occasional cleaning

One aspect not mentioned above, but which to my mind is a substantial advantage of the cup-and-vane type, is that it absolutely need not rely on any measurement technology. Merely by glancing at it you can tell (a) the direction of the wind at any instant, and (b) is it calm, gentle breeze, stiff breeze or hurricane. What more do you need to know? Naturally, you would not expect a technologist to dwell on this aspect...

The principle

But anyway, to the method at hand. When the wind blows through the trees, the tree trunks bend. If the wind blows too hard, the force at the base can pull the roots right out of the ground, but at lesser windspeeds there is a stress exerted on the trunk which is greatest at ground level and is proportional to the wind strength (we'll get to exactly how proportional shortly). You can make a model tree with a plastic rod and an omnidirectional sail on top. Fix the bottom of the rod firmly so that it cannot move, and measure the deflection at or near the base due to the bending moment with a set of orthogonal resistive strain gauges.

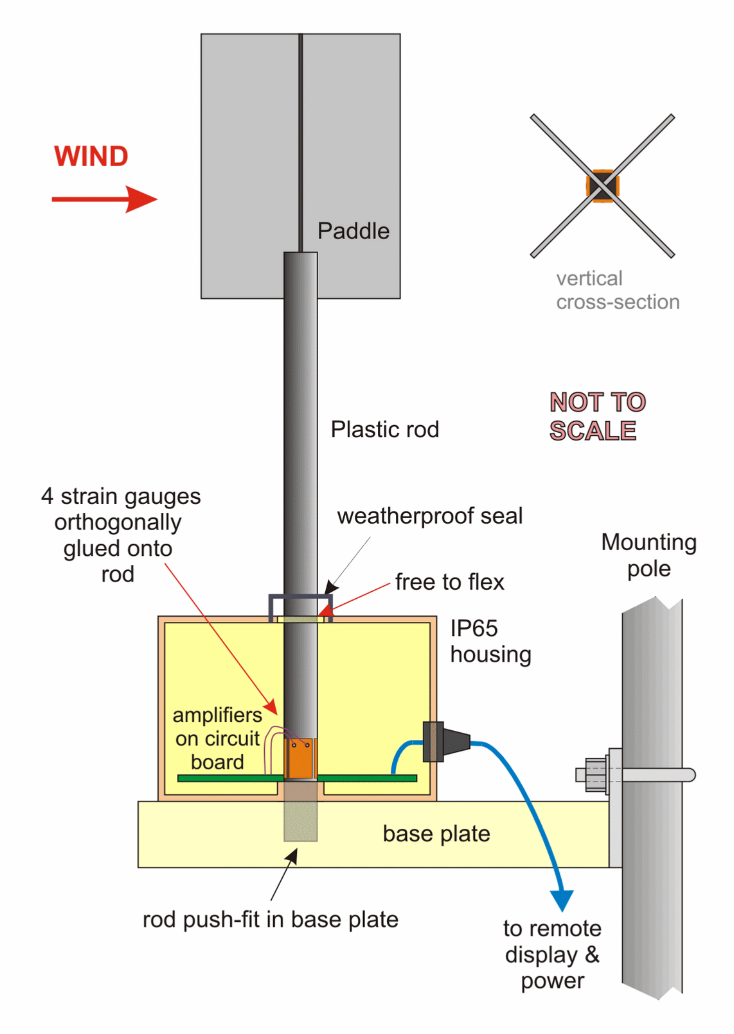

With no wind load, the North-South and East-West pairs should be completely balanced and no voltage is seen at the output of the strain gauge bridge. Any wind deflection will create two voltages from the two N-S and E-W pairs which can be digitised, converted from Cartesian (X-Y) to Polar (r – Θ) coordinates and used to drive a compass rose display.

The strain gauge resistance is proportional to its deflection ("strain") in the longitudinal direction; the specification of the parts used says that they can cope with up to ±2%. Physics theory can give a rough idea of what to expect for a given set of design parameters. It goes as follows:

Wind loading on a flat surface is given [OmniCalculator] as a dynamic pressure:

P = 0.5 x 1.225 x (windspeed m/s)^2

The factor 1.225 kg/m^3 is the air density for standard temperature and pressure. A dynamic pressure P gives a load in Newtons N on an area A m^2 at right angles to the wind direction:

N = P x A

And 1 Newton = 102 grams-force

Working backwards from this, we get the wind speed V in mph (not m/s) for a 10cm x 10cm square "paddle":

V = √(163.26 x N) x 2.24

So for e.g. a load of 0.35N (= 35.7gf) the windspeed is 16.93mph and, importantly, follows a square-root law.

Now we want to know what voltage at the strain gauge amplifier outputs results from a given load at the paddle. The theory of bending [1] says that the strain at one fixed end of a beam from a weight at the other end is related to the flexural modulus of elasticity of the beam material. I'm using 8mm dia PVC rod, for which the *tensile* modulus of elasticity is quoted as 3000 MPa. This may or may not be equivalent to the flexural modulus. But it's easier to forego the maths and just do a measurement. Given a prototype rod with strain gauges attached, and the amplifier electronics, this is no problem.

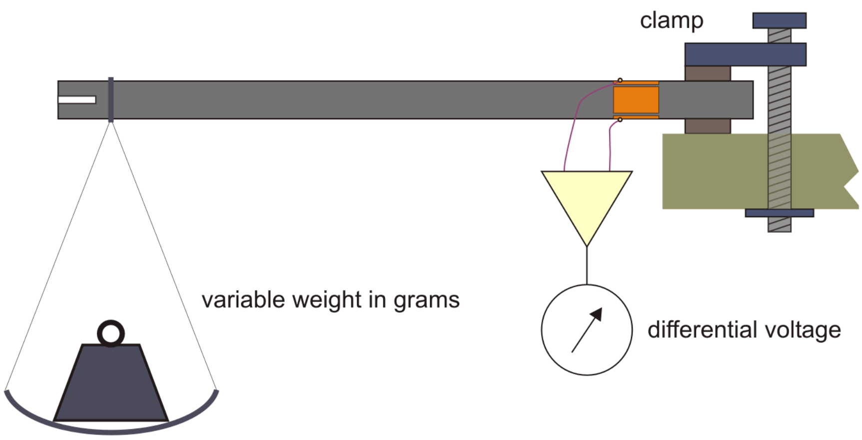

Unsurprisingly, a weight of x grams hung on the end of the rod held horizontal in a clamp, creates a force of x grams-force on the rod. So it's merely a matter of arranging an array of known weights and a sling to hold them in. We can then graph the output voltage for a given amplifier gain directly against the weight and from that, calculate the windspeed to report. There is one slight quirk: the weight of the un-loaded rod of course causes a deflection at the gauges before you start, so you don't start from zero. You can get round this, pretty accurately, by making a second measurement with the rod turned through 180°, and taking the average of the two.

An example result is shown here for a 40cm rod and 10 x 10cm paddle:

[1] *Mechanical Engineering Systems*, Gentle, Edwards & Bolton, Butterworth Heinemann 2001, ISBN 0 7506 5213 6, section 5.4

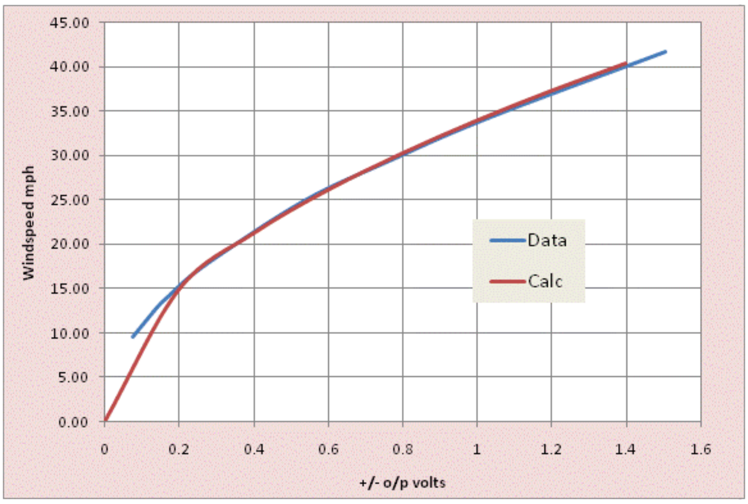

To get the closest approximation using a square-root calculation an offset of 2mph is needed, which has consequences for low values of windspeed (see later); but otherwise the calculated fit is pretty good for a prototype.

System block diagram

The total system comprises the sensor unit, mounted remotely up a pole or on a chimney; this communicates to an in-home Raspberry Pi which hosts a Python software package to display and analyse the data. Everything was designed, built and programmed in-house.

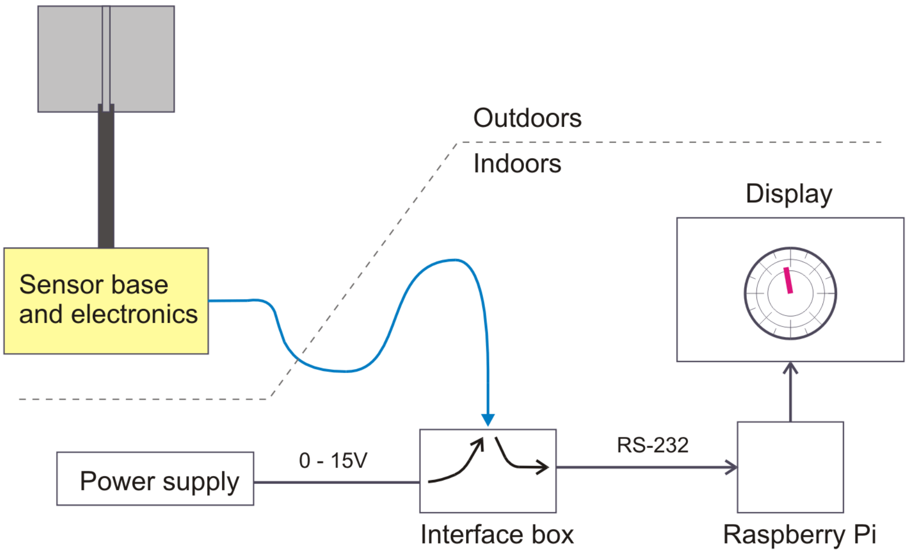

Mechanical design

Strain gauges

For the mechanical part of the realisation, the first question was, where to find the requisite strain gauges. Professional/industrial parts of course exist but cost upwards of £20-£30 a hit. Enter Ebay: for some reason (I'm not speculating) there has been a glut of cheap Chinese-sourced parts for some months, possibly years. An Ebay search reveals numerous vendors of exactly the same part, which is called a “BF350-3AA Precision Strain Gauge”, and which is normally sold in packs of 5 for between £4 and £8, free delivery from China. The device measures 7mm x 5mm on a thin flexible substrate, has a nominal resistance of 350Ω, and can be had with or without thin wires attached. (Other similar parts with different dimensions are available, but this size seems to be very popular.)

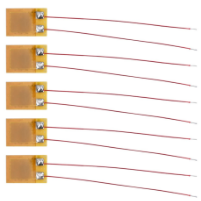

At this price, you can hardly lose, so I ordered a few packs. The first thing to note is that they are incredibly fragile. With a fine-tip soldering iron you can attach the wires yourself, but you can expect to destroy a proportion of units in the process: easier to let the vendor’s experienced operators do it, and find a way to connect to the other end of the wires. (This was a subject of some experimentation before finding a reliable method.)

From a previous project I already had a pack of 8mm dia PVC rod (RS part no 438-6308) and it was merely (ha!) a matter of glueing the gauges orthogonally around the circumference at one end. It was quickly obvious that they needed a flat surface and so a nearly-square cross section was cut 2cm from the bottom end of the rod, which was a bit more than 5mm wide and nearly 10mm long. Geometry says this can be done on an 8mm or larger diameter rod. It would have been nice to have done this on a milling machine, but I can vouch for the fact that it is possible to do it, adequately, with care, using a hacksaw and file.

Glueing such a small part wasn’t that simple; at first I tried cyanoacrylate (Superglue) on the basis that this should cure quickly, but this turned out to be unreliable, as after a short while with even minor flexing the gauge substrate would partially lift from the rod surface. Needless to say the whole of the substrate must remain fully bonded to the rod, and in the end I settled on a miniscule dab of Araldite epoxy, well mixed on a very clean plate. Both surfaces must be cleaned with an IPA/water wash and not touched thereafter. You have to resist the temptation to try it for a good 24 hours after assembly, as it has to cure fully before the strain on the rod is passed to the gauge without slipping and hysteresis. There are undoubtedly more specialized epoxies than plain Araldite which you would use in a professional environment, but so far the Araldite seems to be holding.

Paddle

At the other end of the rod, we need a paddle or sail which is omnidirectional to catch the wind. To simplify construction, a pair of orthogonal aluminium plates, each 14cm H x 10cm W have been used. SSH used a laser cutter to cut a complex paddle shape from plastic sheet; not having access to one, I used a hacksaw and Dremel on 1mm aluminium. And no, I don’t live in a cave.

OK, this is not strictly omnidirectional: the effective cross sectional area varies from unity at face-on at each multiple of 90° to 0.707 at each odd multiple of 45°, and this is without consideration of likely aerodynamic effects. Setting the paddle at 45° to the orientation of the gauges may be a means of mitigating the potential inaccuracy. I haven’t spent the effort to verify this, not having a wind tunnel; a proper design would have to go the distance to be sure. It would be theoretically possible to correct this in the display software, but I’m yet to be convinced that the extra accuracy would be worthwhile.

Cutting a crossed pair of slots in the top of the rod, about 1cm deep, into which the paddle can be shoe-horned, makes for great ease of assembly when you’re at the top of a ladder and perhaps surprisingly has withstood windspeeds up to 40mph without glueing.

Base

It's important that the base of the rod, below the strain gauges, is securely held and is also maintained as near upright as possible; any deviation of the rod from vertical translates to an offset, dependent mostly on the weight of the paddle, which has to be calibrated out electronically. As the intention is to chimney-mount the whole sensor unit via a TV-aerial type clamp, a solid acetal block ended up as being the easiest to machine as well as weatherproof and not too expensive (from The Plastic Shop). The bottom section of the rod with the associated electronics was enclosed in a weatherproof junction box housing which was bolted to this block (see diagram and photo).

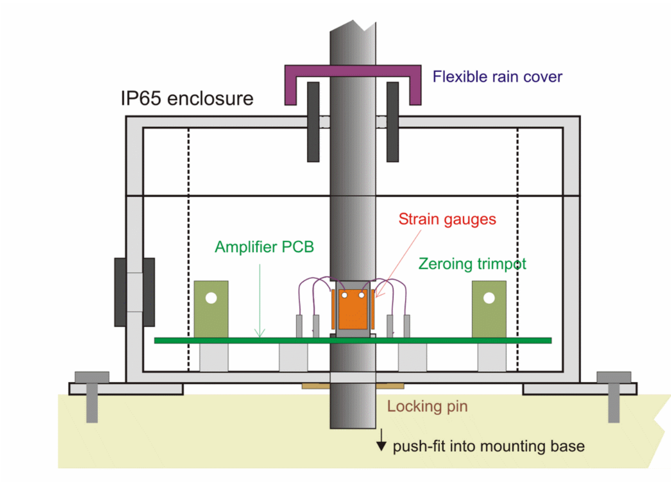

Sensor unit electronics

The electronics at the base of the rod has to:

- Provide a constant current to the 350Ω strain gauge bridge pairs
- Amplify the voltage difference for each pair (North-South and East-West) to provide a voltage against a fixed reference
- Measure the local temperature
- A-D convert the two difference voltages, the reference voltage and the temperature, along with any external signals that may also be wanted
- Sum several measurements of the difference data to smooth out immediate variations
- Transmit the digital data to the in-house system

All this is done inside the housing at the masthead which contains the bottom of the rod with the strain gauges. There is no particularly clever electronics; the two differential amplifiers are INA826 instrumentation amps, which with a 2.5V reference and an LM50B temperature sensor feed the AD inputs of a PIC16F676. A 15V supply feed is stabilised with a 7812 and 78L05 regulator chain, the total supply current being about 40mA, much of which is the drive to the strain gauges. The 0.6W internal dissipation mildly raises the internal temperature a few degrees above ambient and helps (a bit) when the outside temperature is at or near freezing; though this may not really be needed since the ICs themselves are specified down to -20 or -40°C.

The critical components are the two 1k resistors marked with an asterisk in the diagram. These must be well matched, co-located on the PCB and with a low tempco. Any differential drift here will affect the zeroing in nil wind. The value is a compromise between a reasonable current drive from the 2N3702, the common-mode range of the amplifier and the gain required of the amplifier. For the values shown and with a ±0.5Ω swing in the gauges, the differential input to the amplifier is 3mV and its output is 320mV. The 20Ω trimmer takes out the residual unbalance of the gauges and the resistors, and it is the lowest value that can easily be obtained, which is why it is paralleled with even lower (3Ω) resistors to reduce the adjustment sensitivity. A cermet pot rather than a wirewound type is preferred.

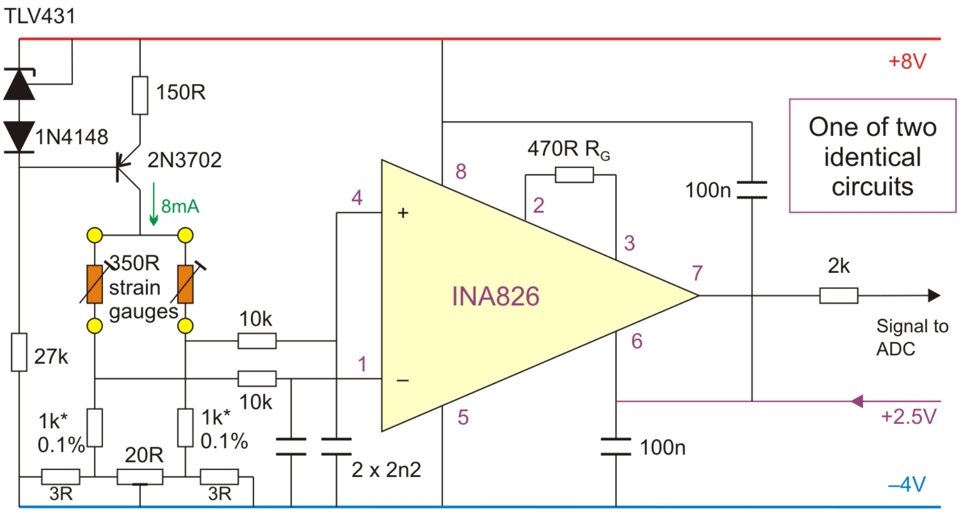

Schematic of one strain gauge amplifier

The PIC software, written in assembler, only has to perform a series of A-D conversions, sum the results and transmit the data at 4800 baud every 2.1 seconds. The timing has been chosen to be a 10-bit conversion every 65ms, 32 of which are summed to create 15-bit data for transmission. Averaging and further signal processing is done at the other end - most of the time the PIC is idling.

Display and recording software

The nice thing about this project is that, as well as mechanical design, electronics design and a bit of physics thrown in, there's a lot of quite interesting software at the display end. which is written in Python to run on a Raspberry Pi.

The data word from the sensor head appears at the RXD (GPIO15) pin on the Pi header every 2.1 seconds and is received using PySerial. There is no serial transmission in the other direction as the Pi has nothing to say to the sensor, any missed or garbled incoming data is simply lost. The idle loop blocks until the serial port indicates a new data packet, then takes the raw data, makes sure its checksum byte is correct and converts the data word into voltages which are put into a queue for processing by the graphics display loop.

Roughly every second, the display loop checks the queue and if there is new data, it takes the N/S and E/W voltages and scales and converts them to polar co-ordinates (R, theta). The last ten values are then displayed as vector lines on a polar graph, so this graph is effectively updated every 2-3 seconds. Separately, the peak and average amplitude values are created and saved every ten minutes; this data can be directly compared to the old-style information from the Davis wind head that is co-located on the same chimney.

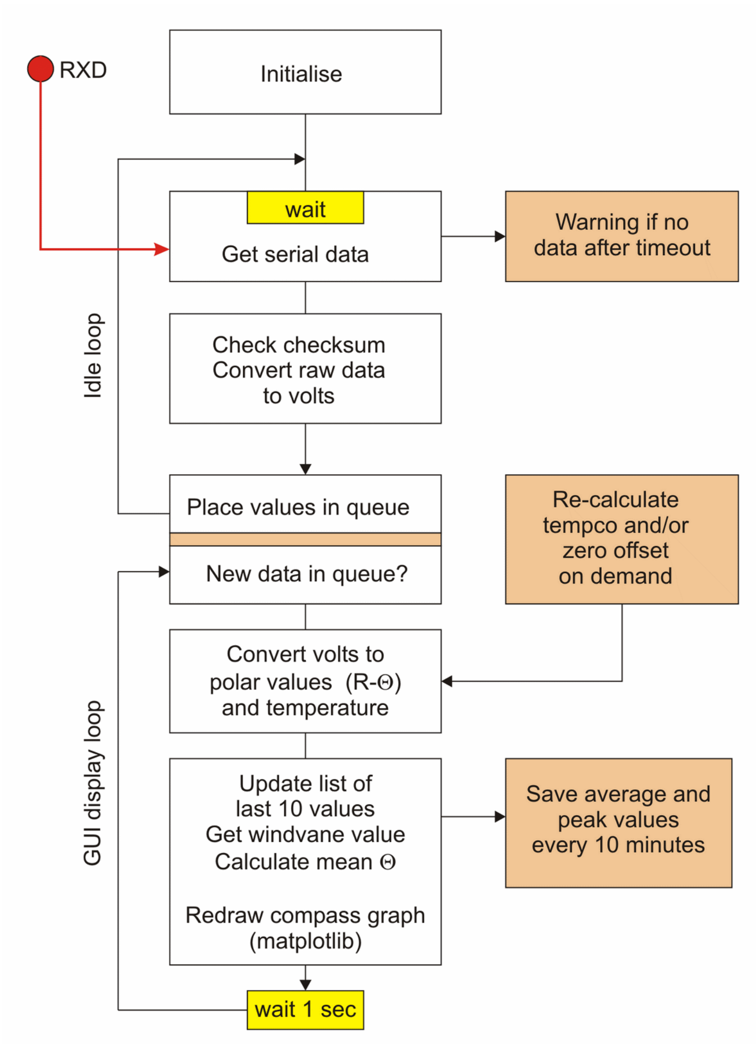

Some subsidiary processing is also done; the Davis windvane value is detected and transmitted in the same data word, and is displayed on the same compass graph. This allows a direct comparison between the sensor's direction measurement, averaged over the last ten values, and the one from the windvane, which is only a few inches away and can be said to see the same wind. The sample display shown below gives an idea of this.

There will be an unavoidable temperature drift (see later). The temperature in the sensor head is included within each transmitted data word and a temperature coefficient correction is applied to the voltages during the pre-display processing; at the same time, a pre-calibrated zero offset is applied to each voltage, and the tempcos and zero offsets can be modified on-the-fly via the GUI in the software if necessary, which it sometimes is.

Polar display

Although the software is doing quite a lot of data processing, the major effort is actually in the drawing of the compass graph on each update. I'd done some previous work for other projects which required graphical display, and had settled on Matplotlib as the mechanism for this. And since I'd discovered that Matplotlib can do pretty much anything in the way of a graphical display it seemed likely that a polar display would not present it any challenge. And this is indeed the case: Matplotlib can create a graph of any type with any format that you can think of. The challenge is working out how to do it. The package has perhaps the most extensive documentation of any I've used, and as a Python library it is of course free to use, but it seems like I've spent weeks going through it to find exactly what parameter has to be set up to do exactly what I want. I doubt very much that mine is the most elegant or optimum implementation, but it does work.

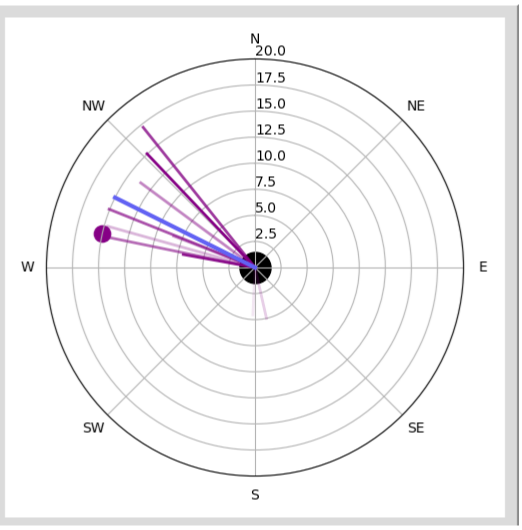

<u>Example display</u>: the blue line is the windvane readout; the purple lines are the last 10 strain gauge unit readings with the oldest ones most faded; the purple dot is the average *angle* of these last ten readings, but *not* related to the windspeed

Development experience

Quite a lot of things have been learnt as a result of this project. Here are the major lessons.

Hyper-sensitivity to low windspeed

Earlier it was mentioned that the windspeed-voltage curve follows a square-root law. The consequence of this is that the maximum sensitivity occurs around the zero point, and so very small amounts of drift, or vibration, or simply mechanical or electrical noise, prevent us ever seeing 0mph when the air is completely still. Effectively, the unit always reads 1-3mph in a random direction in still air. The easiest way to put up with this is to blank out the region of, say, less than 3mph in the compass rose so that no notice is taken of this effect. After all, we're normally more interested in high winds than low. But this is, of course, a disadvantage compared to a cup anemometer, which is quite unequivocal in still air: the cups just don't turn.

Temperature compensation

Of greater concern at low speeds is the problem of temperature drift. The sensors and associated electronics are out in the open and subject to the full range of environmental temperature variation. In this case, I've seen a range of -5 to +45°C logged inside the sensor enclosure. The higher temperature occurs in full sun with a measured outdoor temperature of 33°C, clearly the internal dissipation and solar gain giving a particularly extreme value (for England: this is probably not extreme for, say, Saudi Arabia).

So the ability to cope with a temperature swing of 50°C is desirable. This mandates temperature compensation of the received voltage values.

The sources of temperature drift are:

- Variation of the drive current to each pair of gauges; this is controlled by the TLV31/2N3702 circuit as shown in the schematic, and as long as the gauges are well balanced its effect will be low;
- Variation of the offset voltage in the INA826, this is specified at typically 0.4µV/°C, so a maximum variation of 0.02mV;
- Variation of the 1k thin-film resistors; the ones I used have a quoted tempco of 25ppm/°C; because they are a couple of cm apart on the PCB I would guess there might be a worse-case temperature difference of 4°C which would mean

a resistance difference of 100ppm or 0.01%, implying a worst-case 0.8mV voltage difference;

- Variation of the strain gauge resistance between pairs. It would be preferable (and possible) to match each pair, by hand, accurately to 0.03%, but if you don't do that, their individual tolerance is no better than 0.4Ω or 0.11%. The particular tempco of these gauges isn't quoted, but they use constantan which has a tempco of 0.00001 ohms/ohm-°C, i.e for a 350Ω device the resistance variation over 50°C will be 0.175 ohms or 0.05%. This translates to a 1mV variation. The greater effect will be felt if there is a temperature gradient from one side of the rod to the other, which may happen when the sun shines on just one side of the enclosure;
- Temperature coefficient of expansion of the rod itself, in the region of the strain gauges. This is quoted as $8.\ 10^{-5}$/K or 0.02mm over a 50°C range and a 5mm length (the length of the strain gauge). If the temperature profile from one side of the rod to the other is small, i.e. not in bright sunlight, both gauges in a pair will see the same variation of expansion and the net effect will only be significant if the gauges are not well matched.

Now, a proper professional design practice would be to do a full analysis of all these parameters in the hope that at least some of them could be minimised or shown to be negligible. I haven't done that, but I have seen several mV of variation as the day-night-day temperature varies by 20°C or so, which substantially affects the reading at low windspeeds - potentially reporting 5-8mph when it is virtually calm. As a result I included a temperature compensation routine in the analysis and display software, using the reported temperature from the unit.

This requires that the voltage variation versus temperature for each strain gauge pair is determined separately, particularly as the various contributors may create either a positive or negative coefficient for the pair, and it is by no means certain that the result for the total tempco would be linear. But by applying an individual correction to each voltage before it is converted to an R-θ value this problem seems to be averted. As a result, I've seen the nil-wind value remain below the 3mph threshold mentioned above over a whole day's temperature excursion, which is perhaps acceptable.

Reliability of connections and long-term drift

Nevertheless, re-zeroing the readings when you know that there is nil wind still seems to be needed every few weeks at least. And this leads on to the problem of long-term drift. If we take out temperature variations there are still going to be other factors which disturb the zero reading.

The first version of the design, which isn't discussed here, had the strain gauges on the rod mounted quite separately from the electronics. A 20cm-long cable, each wire separately soldered to one of the strain gauge flying leads, was taken via a weatherproof connector into the electronics box and then to a board connector to the PCB. This was hideously unreliable, for two reasons:

- There was no simple way to make a sufficiently weatherproof and robust connection between the gauges and the cable. Every few days, one or another of the connections, or the gauge itself, would break. I got through a lot of gauges, heatshrink tube and Araldite with this approach.
- Even if a connection didn't break, I soon realised that the connector contact resistance, specified at several milliohms, was in no way benign. Zero jumps of several mph which couldn't be explained by temperature were found to be due to connector uncertainty. Fully-gold-plated contacts would have been better, but having no connectors in the interface was best.

Hence, version 2 of the design which you see here, has dispensed with the connectors by placing the bottom of the rod with its strain gauges inside the electronics housing, and soldering the flying leads directly to pins on the PCB. This removes the unreliability of the connector interface, and it also ensures that all temperature-affected parts of the sensor assembly are in the box, close to where the temperature is measured, making for more effective temperature compensation.

Long-term drift is now much less noticeable; that due to the electronics should be negligible, and it is probably dominated by imperfections in the gauges themselves, and particularly in their fixing to the rod.

Resonances of long rod structures

And then we come to the final nail in the coffin, so to speak.

When you observe trees bending in the wind, you see that they don't bend consistently only in one direction, against the wind. They thrash around, backwards and forwards. The rod with its paddle does the same thing, it oscillates with the slightest wind. Even with a steady breeze it does this; with a gusting wind, it does it more. Not only that, but the oscillation tends to have a sideways component as well. Of course, every motion in whatever direction is picked up by the gauges.

The frequency of the oscillations is determined by the physical parameters of the rod, acting as a (slightly) damped inductor, and the mass of the paddle, acting as a capacitor. Indeed, looking at it as an analogous electrical LC circuit, it's no surprise that it should oscillate at its resonant frequency.

What this means is that individual strain gauge measurements on their own are useless. If nothing can be done to damp the mechanical oscillations[2], then the readings have to be damped electronically, by averaging a number of them, in the hope that the mean of the readings will still indicate the mean direction and speed of the wind. This is done in the sensor's PIC processor by taking a measurement every 65ms, before the summed data is transmitted every 2.1 seconds, giving a total of 32 measurements per "reading".

To some extent, this does indeed work; but only partially. There is still a large amount of fluctuation in the indicated 2-second values. I haven't determined the natural frequency of the rod/paddle combination and this is work yet to be done, in the hope that a tweak to the sampling rate might improve matters. On the other hand, it seems likely that this is a fundamental limitation of the method and it will always create a degradation in the indicated result.

Accuracy of peak windspeeds

One discovery that is positive is that although the reported wind *direction* every 2.1 seconds is somewhat arbitrary, as the windspeed increases it becomes more correct. A comparison of the reported direction with the measured vane direction on a per-minute basis has shown something of a correlation, although not strong, with the peak speed; that is, as the speed increases, the average difference between the vane direction and the reported sensor direction goes down. So that, if you want to know an accurate direction, it's helpful to have strong winds.

The correspondence between peak windspeeds shown by this unit, and the Davis cup anemometer, on a 10-minute basis is much better. Here is a 14-day graph of the wind sensor logged peak windspeed versus the Davis anemometer ditto.

[2] Something probably could be done, but it would be mechanically ugly

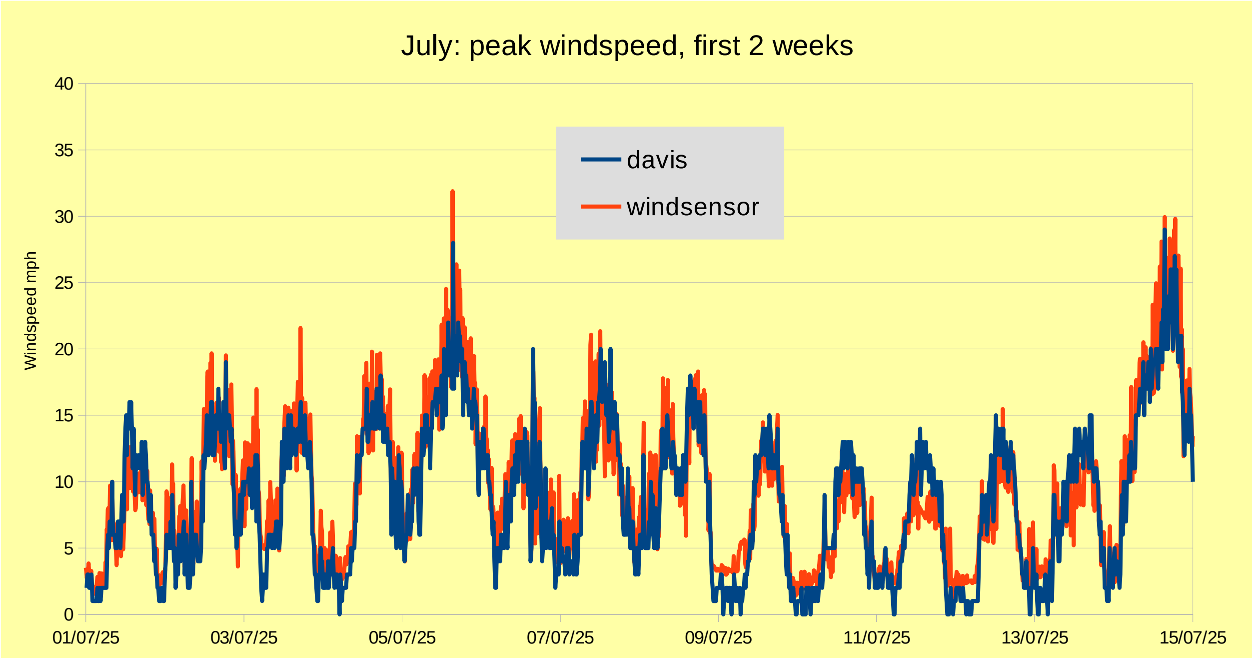

Conclusion

The work detailed here has shown that the principle of strain gauge wind measurement can produce useful results, but this mechanical design is probably not the best way to do it. There are two particular limiting factors which work against it:

- The sensitivity to electrical drift and very small movements of the assembly in nil or light winds, which translates to random direction indication at low speed;
- The resonant characteristic of the rod/paddle assembly, such that the oscillations in the individual measurements must be smoothed out by software averaging.

Also, the installation must make sure that the assembly is carefully levelled to minimize the bias that results from the weight of the paddle. On the other hand, the peak amplitude indication at higher windspeeds is quite accurate and if this were the only requirement, the method may well have its advantages.

SSH’s mechanical design does away with the long rod and has a better paddle design, and this would quite probably deal with the major disadvantages, so if I was doing it again I’d look much more closely at that. It would also be possible to substantially reduce the current drain and transmit the data over a wireless link, with a solar cell/battery power source so that there is no need for a wired connection. That wasn’t in my original brief; but it could well provide next year’s challenge!

Return to Tim’s Notes