CIRCUIT CELLAR®
THE MAGAZINE FOR COMPUTER APPLICATIONS

Circuit Cellar, the Magazine for Computer Applications. Reprinted by permission. For subscription information, call (860) 875-2199, or www.circuitcellar.com. Entire contents copyright ©2006 Circuit Cellar Inc. All rights reserved.

FEATURE ARTICLE

by Ingo Cyliax

Internet-Connected Sonic Anemometer

Ingo's Internet-based sonic anemometer gauges wind velocities by measuring the speed of sound in air. Read on to learn how the system works. Ingo has included all the information you need to get started on an anemometer of your own.

I recently became interested in anemometers when I took a new job in the weather instrument industry. One of the most common instruments used in weather stations to measure wind velocity is the anemometer. Cup anemometers feature several half-spherical, hollow cups attached to radial spokes (see Figure 1). It's intuitive to see that the coefficient of drag for a spherical cup differs depending on where the air is pushing. Basically, the drag is high for the hollow side in comparison to the closed side.

This difference in drag generates torque that spins the entire assembly. A tachometer then measures the resulting rotational speed, which has a linear relationship to the wind speed. Typically, the tachometer is either implemented as a small generator that generates a voltage into a load or a Hall effect or optical sensor that generates pulses relative to the rotational velocity.

Another common method is to use a small wind turbine. The wind drives a propeller or turbine wheel. The resulting rotational speed is measured to give an indication of the wind speed. Using cup and turbine anemometers can be disadvantageous because they require a significant amount of minimum wind to overcome friction and drag. They have momentum, so they don't respond well to wind gusts and tend to average out the speeds.

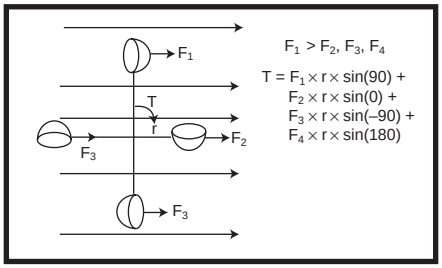

Figure 1—*The rotational speed of a traditional cup anemometer is determined by the difference in drag force and indicates wind speed. A separate vane is used to indicate the wind's direction.*

There are other instruments for measuring wind speed. Pitot tubes are used in applications where the wind comes from a specific direction (see Figure 2). For instance, when measuring a plane's airspeed, the wind always comes from the front of the plane with only small angles of deviation. A pitot tube works by measuring the dynamic pressure of the wind as it enters a small hole pointed toward the wind. The tube compares the dynamic pressure to the ambient air pressure. The dynamic pressure is easy to calculate:

$$P_{DYN} = P_{STAT} + 0.5 \times (rho) \times V^2$$

where *rho* is the air's density. As you can see, you have to subtract the static pressure in order to compute the speed:

$$V = \sqrt{\frac{P_{DYN} - P_{STAT}}{0.5}}$$

A disadvantage is that a pitot tube anemometer is accurate only if the wind blows directly into the pressure port.

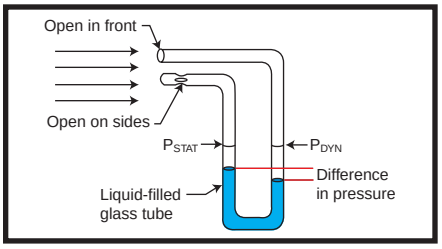

Figure 2—*A pitot tube anemometer is common on airplanes and in applications where the wind comes from the same direction.*

A hot-wire anemometer measures the heat carried away by a moving air mass (see Figure 3). The higher the wind speed, the more heat carried away (i.e., it cools the wire).

Pitot tubes and hot-wire anemometers are advantageous because they don't have moving parts that need to be sealed from the elements. Anemometers are often installed in remote weather stations where they're hard to maintain, so having one with no moving parts is a perk.

Using a Doppler lidar is an interesting way to measure the effects of wind speed. A laser beam is sent out and particles in the air (aerosols or water) scatter it. Some of the scattered beam is reflected to a sensor that measures the shift in the laser's wavelength when the aerosols are moving. You can measure their radial speed to the laser by measuring the laser's Doppler shift (or an RF carrier modulated on the laser).

A fringe-type laser anemometer projects a field of interference fringes and detects how aersols are illuminated by the fringe pattern as it moves through the field. complex devices require powerful lasers to get any range. They require high-frequency, fast electronics to make them work with any resolution or accuracy. They are usually limited to scientific applications. For example, in systems that measure wind speed and aerosols

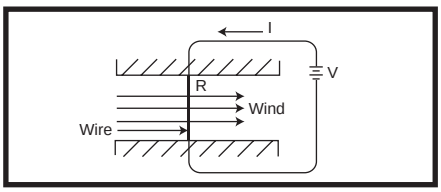

Figure 3—*A hot-wire anemometer measures the wind by how much it cools the resistively heated wire.*

tens of kilometers in the atmosphere, or measure wind speeds in wind tunnels.

A sonic anemometer measures the speed of sound in air. In this article, I'll explain how you can measure wind speed with one. At this point, you might be thinking, "So what!" The perceived speed of sound in air is actually the speed of sound in static air plus or minus the speed of the wind (see Figure 4). In other words, the wind speed is the difference between the measured speed of sound and the actual speed of sound in air.

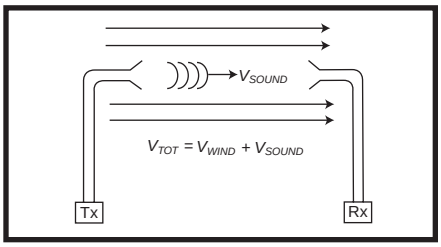

Figure 4—*In a sonic anemometer, the wind speed adds to or subtracts from the speed of sound in air depending on the sound's direction.*

SPEED OF SOUND IN AIR

The speed of sound in air depends mostly on the temperature. It depends slightly on humidity and the air's actual mixture of gases. Luckily, you have to worry only about the temperature. Furthermore, if you're only interested in the wind speed, you can just factor out these effects.

The standard speed of sound is usually determined by one of the following:

$$C_{AIR} = 331.5 + 0.6 \times T_C$$

or

$$C_{AIR} = 20 \times \sqrt{T_K}$$

These are approximations, and they don't account for humidity. A more precise instrument would use more elaborate models for the speed of sound in air.

How do you go about measuring the speed of sound? The most reliable methods involve ultrasound range measuring. Ultrasonic transducers, especially for the 40-kHz band, are readily available. I've seen them used in numerous robotics projects.

In an ultrasonic ranger, a short burst of ultrasound (8 to 16 cycles) is transmitted via a transducer. Another transducer then picks up the reflections of the sound. The delay between the transmitted burst and the received burst is called time of flight (TOF), which can be converted to distance if you know the speed of sound.

Photo 1—*I built my three-point anemometer with PVC pipe. The distance between the adjacent sensors is 0.36 m.*

$$\text{Range} = T_{TOF} \times \frac{C_{AIR}}{2}$$

Although ultrasonic transducers are bidirectional devices (i.e., they function as transmit and receive transducers), they usually use a separate transmitter and receiver transducer. The transmit transducer has a tendency to ring for some time after the excitation waveform stops. This limits the short range of such a system.

Transducers can be optimized for transmitting or receiving. If you substitute one for the other, you'll get reduced performance. Circuitry for rangers using separate transducers is available in many robotics publications and datasheets for transducers.

Measuring the actual speed of sound is easy if you have a calibrated distance. In sonic anemometers, the transducers are mounted in a fixture that fixes a direct distance between transducers. If you measure the direct time of flight between the transducers, you can compute the perceived speed.

$$C_{PERCEIVED} = \frac{\text{Distance}}{TOF}$$

This speed is the speed of sound in the air plus the speed the wind exerts.

$$C_{PERCEIVED} = C_{AIR} + C_{WIND}$$

If you determine the air temperature with a thermometer, you can compute the wind speed simply by applying one of the earlier expressions:

$$C_{WIND} = C_{PERCEIVED} - (331.5 + T \times 0.6)$$

Now, here comes the clever part. If you measure the speed in both directions

along the same path, you don't need to know the temperature.

$$C_{PERCEIVED}12 = C_{AIR} + C_{WIND}$$
$$C_{PERCEIVED}21 = C_{AIR} - C_{WIND}$$

By combining a the equations, you get:

$$C_{AIR} = \frac{C_{PERCEIVED}12 + C_{PERCEIVED}21}{2}$$
$$C_{WIND} = \frac{C_{PERCEIVED}12 - C_{PERCEIVED}21}{2}$$

You can also derive the current temperature:

$$T = \frac{C_{AIR} - 331.5}{0.6}$$

Of course, the system needs to transmit and receive in both directions along the path. This requires the transceiving transducers or a set of transmit/receive transducers.

Ringing in the transmit transducer isn't a big problem because the turnaround time is slow. The system waits for the burst to travel the distance to the other side before it switches direction. But the ring has a tendency to stretch the pulse by adding more cycles to it. This seems to work. Well, not quite. It only determines the wind speed along the path. When the wind blows from an angle, you can see only a portion of the wind speed. When the wind is orthogonal to the path, you can't see any wind speed at all.

To fix this, you can configure the anemometer to measure in different paths. The cross (orthogonal) is a typical configuration. Use the same technique to measure the wind speed along each of the two paths (*x* and *y*). To compute the absolute wind speed, simply add the squares of each component and take the square root (magnitude equation):

$$C_{WIND} = \sqrt{C_{WINDX}^{2} + C_{WINDY}^{2}}$$

You can also derive the wind direction:

$$\text{Direction} = \text{Atan}\left(\frac{C_{WINDY}}{C_{WINDX}}\right)$$

Other configurations are also possible. I used a triangle (60°).

You might have noticed that you aren't measuring the Doppler shift of the sound burst. A common misconception about using a sonic anemometer is that you can measure the wind speed by measuring the Doppler frequency shift of the ultrasonic signal.

Recall the high school physics experiment relating to train whistles changing pitch as they approach and recede from a stationary listener. Of course this is true, but in the sonic anemometer setup, both the listener and the train whistle appear to be moving at the same speed. So, it's more like listening to the train whistle while you're on a car in the back of the train (i.e., the pitch doesn't change) even though the train may be moving at great speed.

The Doppler signal in a sonic anemometer measures the change in wind speed, which can give an indication of the wind speed's stability at the moment it's measured. You can use this to measure turbulence and fluctuations in wind speed. Of course, this involves being able to measure the changes in the received signal's frequency. You can use a fast Fourier transform (FFT) to look at the purity of the spectrum. You may also need a transducer with a wide frequency range. Many ultrasonic transducers are tuned for a specific frequency and attenuate if the received sound frequency is out of the pass-band. Check the transducer's datasheet.

IMPLEMENTATION

I built a three-transducer, three-path sonic anemometer. Each transducer can see the other two transducers by arranging them in a triangle (60° angles). Three paths is a good compromise between minimizing the number of transducer channels and maximizing the number of paths for measuring the wind speed. The aforementioned two-path, four-transducer anemometer needs one of the transducers.

The transducer I selected is optimized for transceiving. It has a beam angle wide enough to see the others with a 3-dB signal reduction. For transmitting an AC signal, up to 20 V_{PP} is applied to the terminals. When the transducer is connected to a high-impedance input, it behaves like a microphone and generates a voltage in the load.

The transducer's design dictates the resonant frequency (40 kHz) and beam angle. I could have used a different frequency transducer (125 or 250 kHz), but there is less absorption loss in the air at a lower frequency. Furthermore, lower frequencies make the processing requirements less demanding.

The transducers are mounted in a PVC pipe structure (see Photo 1). PVC

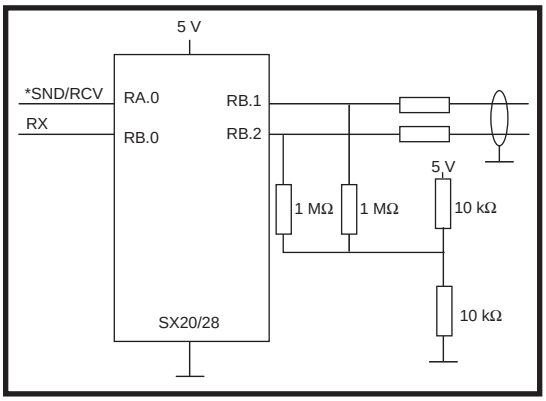

Figure 5—*The transmit/receive circuit is based on a Scenix SX20/28 processor. The transducer is attached to the pins that serve as both an analog comparator for receiving and digital I/O, which drive the transducer in a push-pull configuration. I used two-conductor microphone cable to reduce crosstalk.*

tube is great for building these structures because it's easy to work with and a variety of fittings are available. The transducer fits in the tip of each arm and is held in place with a little RTV. I originally used twisted-pair cable to connect to the transducers, but that was a mistake. I've since switched the wiring to two conductors with shield microphone wire. The two conductors are isolated from the ground-shielding wire. The shield is terminated on the electronics end. This reduces crosstalk between the transmit and receive wires and it eliminates ground-loop oscillations.

The front-end electronics were the most difficult to come up with. I had to switch the transmission circuitry, which is relatively low impedance with the high-impedance inputs to the receiver. After trying a different scheme, I finally settled on a solution featuring a Scenix SX28/20 microcontroller (see Figure 5). The SX28/20 microcontroller has attractive features like symmetric CMOS outputs, which can source and sink up to 30 mA each. I used the CMOS outputs as push-pull drivers and a supply voltage of 5 V to drive the transducer with a 10-V voltage differential with plenty of current.

I can switch off or tristate the outputs to put them into a high impedance state. A built-in analog comparator has a low offset voltage (10 to 20 mV) and short delay (25 ns) that make it suitable for a sensitive zero-crossing detector. The processor can read the comparator's output directly. The output can be routed to a separate output pin. The input pins can be shared with the output pins.

In order for the received signals to be in the comparator's voltage range, a simple resistor network must bias the inputs to a point approximately halfway between V_{CC} and GND when the pins are in a tristate condition. It turns out that the 10- to 20-mV offset voltage is perfect for this application. Under normal conditions, the received signal has up to approximately 100 mV for the selected range and driver voltage. That's enough to give a reliable zero-crossing detection, but isn't too sensitive to pick up noise transients.

The front end has two interface signals for control. When the *SND/RCV pin is pulled low, the front end enters Transmit mode. A 16-cycle, 40-kHz square wave burst is then sent to the transducer. When the *SND/RCV pin is pulled up, it switches to Receive mode and sends the comparator's output back via the Rx line.

The SX28/20 microcontroller is available in a 20-pin SSOP package, so you can integrate the entire front end in a package that can be soldered directly to the transducer and fit inside the PVC tube that holds it. This approach eliminates the need for high-impedance wires in the anemometer. Four wires (5V, *SND/RCV, RX, and GND) can be routed via twisted-pair wire to the sensor and front-end electronics with an SBC (see Figure 6). If RS-485 drivers and receivers are used, it's possible to route the signals for long distances.

I used a Rabbit Semiconductor RCM3000 development board for the front-end electronic interface. The processor has a high-speed pulse capture that enables me to take precise measurements of the time of flight. In addition, it has an Ethernet interface and an optional Wi-Fi connection kit. The prototype board has a regulator that accepts 9 to 12 VDC, which enables me to connect it to a sealed lead acid battery in case I want to run the system disconnected from the main.

For portability, economical solar panels and a charger are available for 12-V systems. With a Wi-Fi connection, I can use my laptop to connect to the instrument and communicate with it using a 'Net interface or FTP to upload wind data.

In addition to the aforementioned features, the software library comes with fairly fast floating-point routines. The TCP/IP stack and web server is also included.

The SBC continuously goes through an acquisition cycle to measure the TOF delays between each transducer. For each ping, you can collect two paths by using the two pulse capture registers. So, with a round-robin cycle of three pings, you can collect six path times. There are forward and backward times for each of the three physical paths (T_{12}, T_{21}, T_{13}, T_{31}, T_{23}, and T_{32}). The path lengths for the three segments are measured (D_{12}, D_{23}, and D_{31}) and adjusted for processing delays.

The pulse capture happens to be in the hardware, so there aren't any real time constraints on the system other than the transition from sending the transmission pulse and the arming of the capture time register. But this is a short time that can be handled by simply turning off the interrupts for a few instructions. It doesn't affect the performance or add delays in processing when dealing with the TCP/IP protocol stack and web server.

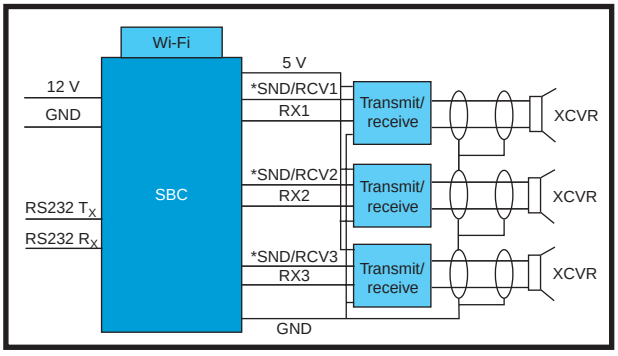

Figure 6—*The system is divided into three subsystems. The SBC handles all of the communications. The transmit/receive circuit interfaces and modulates the transducers. The transducers are in the tips of the actual sonic anemometer assembly.*

To measure wind speed with resolution of 1 meter per second, I need a timing resolution of 1 part in 300. The path is 0.36 m in my instrument, which results in a 1.2-ms TOF delay. A 1/300 resolution implies a timing resolution on the order of 6 µs, which the hardware capture in the RCM3000 has no problems with.

The system uses a different geometry than a traditional cross anemometer using orthogonal paths. The math is a bit more involved, but you can add all of the vector components and use trigonometry.

Assume that the transducer is anchored at the origin and two and three are anchored at 0° and 60° angles (see Figure 7). You must first compute each wind vector:

$$W_{12} = \frac{\frac{D_{12}}{T_{12}} - \frac{D_{21}}{T_{21}}}{2}$$

$$W_{23} = \frac{\frac{D_{23}}{T_{23}} - \frac{D_{32}}{T_{32}}}{2}$$

$$W_{13} = \frac{\frac{D_{13}}{T_{13}} - \frac{D_{31}}{T_{31}}}{2}$$

Following this, add all of the vector components:

$$W_Y = W_{12} \times \sin(60) + W_{23} \times \sin(-60)$$

$$W_X = W_{12} \times \cos(60) + W_{23} \times \cos(-60) + W_{13}$$

Lastly, compute the absolute wind and angle.

$$\text{Wind} = \sqrt{W_Y^2 + W_X^2}$$

$$\text{Angle} = \text{Atan}\left(\frac{W_Y}{W_X}\right)$$

ADDITIONAL IDEAS

As I mentioned earlier, you can extract the Doppler information to compute variations in the wind speed. For most applications, however, it's enough to take samples and compute a running average and variance to get an idea about the wind's stability. Because the TOF distance is short and the sensitivity to echoes and reverberation is low, a firing time of 10 ms is easy to achieve. This results in a cycle time of 33 Hz.

My anemometer computes only a two-dimensional wind profile. But by adding an extra transducer to the vertical plane, I can add three more paths to the instrument and measure a three-dimensional wind profile.

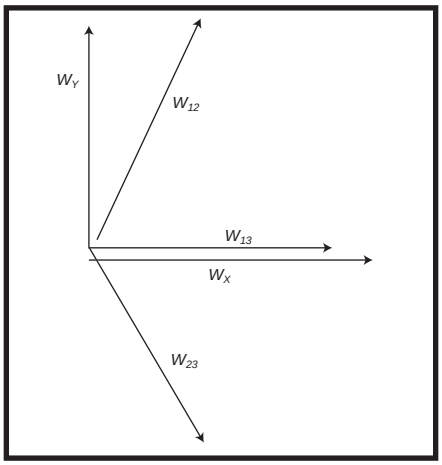

Figure 7—*Computing the vectors in a three-point anemometer is a bit more involved, but it's still manageable.*

Well, this project has been a lot of fun. For my next project, I'll try to miniaturize the front electronics and attach them to the transducers. I'll also weatherize the system so I can use it outside for extended periods of time. A friend of mine is interested in placing a stand-alone unit with a GSM modem near his favorite surfing spots in order to obtain wind conditions. Surf's up!

Ingo Cyliax is the senior research engineer at All Weather, Inc., a Sacramento-based company specializing in weather instruments and systems. He has experience in many areas, including hardware and software design, and has worked at Z-World/Rabbit Semiconductor, Purdue University, Indiana University, and ETH-Zurich. When Ingo is away from the bench, he enjoys hiking, camping, or simply sipping coffee at his favorite sidewalk cafe. You may contact him cyliax@ezcomm.com.

SOURCES

RCM3000 development board
Rabbit Semiconductor
www.rabbitsemiconductor.com

SX28/20 Microcontroller
Ubicom, Inc. (formerly Scenix)
www.ubicom.com