HA&BC

CIRCUIT CELLAR INK®

THE COMPUTER APPLICATIONS JOURNAL

#68 MARCH 1996

SPECIAL BONUS SECTION: HOME AUTOMATION & BUILDING CONTROL

INDUSTRIAL DESIGN

Lightning Protection Techniques

Thick-Film Hybrid Technology

Microvolt Data Acquisition

HA&BC: Wind Speed and Direction Interface for the HCS II

$3.95 U.S.
$4.95 Canada

TASK MANAGER

The Mensa House?

Say you want to sell your house, but you've spent considerable time and money automating it. You've installed wiring and jacks in every room; cameras at the front door and in the garage; speakers in the walls of the living room, kitchen, and family room; infrared repeaters and motion detectors throughout; and a host of other permanent improvements.

How do you give potential buyers an idea of just how "smart" your house is? Just assign it an IQ!

Sound silly? Well, it's a problem builders and retrofitters alike are beginning to face in the home market. Consumers, appraisers, and lenders have no way to assess the value of such features. With more and more houses now being prewired at the time of construction, the problem is only going to get worse.

So, David Gaddis and the folks at the Home Team have initiated the Home IQ Development Program. They have put together a group of industry insiders to hammer out a standard method for measuring the amount of intelligence built into residential and commercial buildings. The effort is just getting underway, but they've made it easy for you to follow. Just connect to their Web site (**http://www.hometeam.com/**) for more information.

THIS MONTH

Speaking of smart homes, we kick off this installment of *Home Automation & Building Control* with an article detailing how to add a wind direction and speed interface to the HCS II. Forecasting the weather is still up to you, though.

In the other feature, we explore some basics in designing a home automation user interface. The key is making it user seductive.

Turning our focus to the Industrial Design section, we start with an article on measuring fluid level using capacitive techniques. Next, we address an issue near to many industrial designers' hearts: lightning protection. And speaking of ruggedizing circuitry, we follow with a look at what's new in the area of hybrid circuits.

Shifting from megavolts to microvolts, the next article gives some guidance to designing circuits that deal with microvolt-level signals. Finally, we present an affordable aid for those who can't see indicator lights—a light probe for the blind.

In our columns, Ed relates some real-world development horror stories, Jeff uses stepper motors in an unconventional application, and Tom looks at the latest in the PLD design.

editor@circellar.com

CIRCUIT CELLAR INK®

THE COMPUTER APPLICATIONS JOURNAL

FOUNDER/EDITORIAL DIRECTOR
Steve Ciarcia

EDITOR-IN-CHIEF
Ken Davidson

SENIOR TECHNICAL EDITOR
Janice Marinelli

TECHNICAL EDITOR
Carole Boster

ENGINEERING STAFF
Jeff Bachiochi & Ed Nisley

WEST COAST EDITOR
Tom Cantrell

CONTRIBUTING EDITORS
Rick Lehrbaum
Russ Reiss

NEW PRODUCTS EDITOR
Harv Weiner

ART DIRECTOR
Lisa Ferry

PRODUCTION STAFF
John Gorsky
James Soussounis

CONTRIBUTORS:
Jon Elson
Tim McDonough
Frank Kuechmann
Pellervo Kaskinen

PUBLISHER
Daniel Rodrigues

PUBLISHER'S ASSISTANT
Sue Hodge

CIRCULATION MANAGER
Rose Mansella

CIRCULATION ASSISTANT
Barbara Maleski

CIRCULATION CONSULTANT
Gregory Spitzfaden

BUSINESS MANAGER
Jeannette Walters

ADVERTISING COORDINATOR
Dan Gorsky

CIRCUIT CELLAR INK®, THE COMPUTER APPLICATIONS JOURNAL (ISSN 0896-8985) is published monthly by Circuit Cellar Incorporated, 4 Park Street, Suite 20, Vernon, CT 06066 (860) 875-2751. Second class postage paid at Vernon, CT and additional offices. One-year (12 issues) subscription rate U.S.A. and possessions $21.95, Canada/Mexico $31.95, all other countries $49.95. All subscription orders payable in U.S. funds only, via international postal money order or check drawn on U.S. bank. Direct subscription orders and subscription related questions to Circuit Cellar INK Subscriptions, P.O. Box 698, Holmes, PA 19043-9613 or call (800) 269-6301.
POSTMASTER: Please send address changes to Circuit Cellar INK, Circulation Dept., P.O. Box 698, Holmes, PA 19043-9613.

Cover photography by Barbara Swenson
PRINTED IN THE UNITED STATES

For information on authorized reprints of articles, contact Jeannette Walters (860) 875-2199.

Circuit Cellar BBS—24 Hrs. 300/1200/2400/9600/14.4k bps, 8 bits, no parity, 1 stop bit, (860) 871-1988; 2400/9600 bps Courier HST, (860) 871-0549. World Wide Web: http://www.circellar.com/

READER I/O

SAFETY ALERT!

"Push-Pull Switching Regulator Design and Application" (*INK* 66) has serious safety problems. Individuals building it could seriously injure themselves or property.

The schematic shows the AC input line connected to circuit ground. While this is standard practice in an off-line switcher, it is dangerous to connect a test instrument, such as an oscilloscope, to the primary side of the circuit (everything left of the transformer). You can get an electric shock as the AC line shorts to ground.

The circuit shows that the +5-V output is connected to the UC2525A PMW integrated circuit. This connection ties the output to the primary ground, which is hot in relation to ground.

Also, for the circuit to work as depicted, the +5 Ret must be connected to the primary side ground. This connection eliminates all isolation between the AC line and the power-supply outputs. There are stringent requirements for isolating the primary and secondary sides of a supply. This method is not acceptable. Feedback from the secondary side to primary side is accomplished through isolating devices such as optoisolators.

The design of off-line switching power supplies is complex and there are serious safety issues. Schram's statement, "In other words, about 0.5 A is drawn from the source-a handy number if you should do something silly like fuse the input" would certainly not be approved by any national or international safety agencies. There always needs to be some sort of limiting protection for off-line switching power supply to be approved.

Schram's supply does not provide for overload or short-circuit protection. Even a momentary short circuit on an output would cause one or both of the power FETs to fail in a shorted condition. Users would certainly wish they had fused the input line if this occurred.

I believe you have done a disservice to your readers by publishing an article like this that oversimplifies the design, construction, and safety issues relating to off-line switching power supplies.

Howard G. Corlett
San Ramon, CA

KEEP IT SIMPLE

I enjoy reading your magazine, but I'm a bit put out by all the complications people seem to require now. I found this particularly true with "Microprocessor-Controlled High-Voltage Power Supply" (INK 66).

The article inspired me to consider the many other ways I'd do it! The supply presented doesn't even have good regulation or response time.

However, my main reason for concern is that there may be a safety issue. Though the author may have realized and tested for it, he didn't address it in the article.

In the last paragraph on page 23, he says,

> The microwave-oven transformer I used has one of its secondary leads cut short and connected to the core ground. I removed this lead from ground..

What is missing is a high-potential test for the transformer secondary to the core after the core connected lead is floated.

A designer and winder of HV transformers, who knows a lead is going to the core, could easily not put sufficient insulation between the core and winding, especially with a microwave-oven component. Such parts are usually cheap. It's nearly impossible to insulate the full HV potential on both ends of the winding. I've seen a lot of these transformers, and they don't look great, though I've not high-pot tested to core on any.

There's still lots of room for straightforward analog and simple digital design which doesn't require large development support.

S. **Premena**
Boulder, CO

Contacting Circuit Cellar

We at Circuit Cellar Mencourage communication between our readers and our staff, so have made every effort to make contacting us easy. We prefer electronic communications, but feel free to use any of the following:

Mail: Letters to the Editor may be sent to: Editor, Circuit Cellar INK, 4 Park St., Vernon, CT 06066.
Phone: Direct all subscription inquiries to (800) 269-6301.
Contact our editorial off ices at (860) 875-2199.
Fax: All faxes may be sent to (860) 872-2204.
BBS: All of our editors and regular authors frequent the Circuit Cellar BBS and are available to answer questions. Call (860) 871-1988 with your modem (300–14.4k bps, 8N1).
Internet: Letters to the editor may be sent to **editor@circellar.corn.** Send new subscription orders, renewals, and address changes to **subscribe@circellar.com**. Be sure to include your complete mailing address and return E-mail address in all correspondence. Author E-mail addresses (when available) may be found at the end of each article. For more information, send E-mail to **info@circellar.com**.
WWW: Point your browser to **http://www.circellar.com/**.
FTP: Files are available at **ftp://ftp.circellar.com/pub/circellar/**.

NEW PRODUCT NEWS

Edited by Harv Weiner

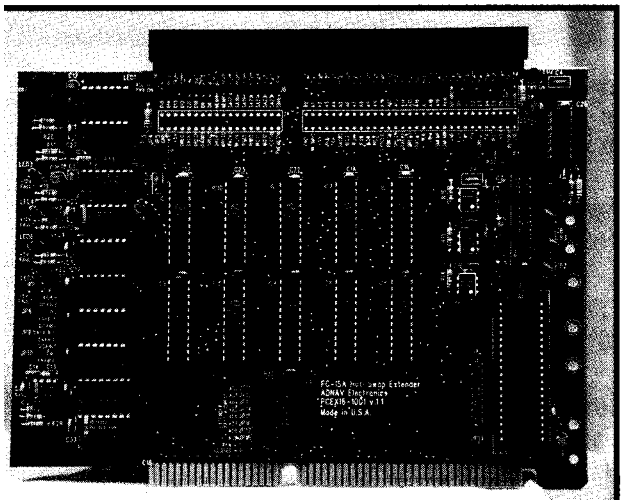

HOT-SWAP EXTENDER CARD

ADNAV Electronics announces the **PCEX16-1001**, an active hot-swap extender card for the PC ISA bus. The user can insert 8- or 16-bit PC add-in cards under test into the ISA bus and remove them without turning the computer off and on. The wait for computer bootup and initialization is eliminated.

The hot-swap extender card uses high-speed solid-state switches to connect the ISA bus's digital lines to the add-in card under test with negligible signal delay. Power lines are also switched. Specialized self-resetting circuit breakers automatically protect the computer and add-in card from power-bus short circuits. The extender's four-layer PCB design leads to low-noise operation.

With the extender, hot swapping is as simple as pressing a push-button switch to power the add-in card down and up again without crashing the computer. Timing logic in the extender card connects bus signals in a sequence that ensures the add-in card can restart its activity after hot swapping. Alternatively, the extender may be controlled through software, simplifying the design of automated test systems.

The extender offers a number of features for testing and diagnosing add-in cards. It clearly labels test points to probe all add-in ISA bus signals as well as the controls that allow selective disabling of data bus, IRQ operations, and DMA transactions to troubleshoot logic collision problems. In addition, LEDs indicate over-current fault conditions on ±5- and ±12-V power buses.

The PCEX16-1001 Hot Swap Extender card sells for $295. A kit version is available for $225.

ADNAV Electronics
58 Chicory Ct. • Lake Jackson, TX 77566
(409) 292-0988

#500

SINGLE-WIDTH VIDEO MODULE

Ariel announces the VPT-40, a video I/O module for Texas Instruments' TMS320C4*x* family of digital signal processors. Packaged in a single-width TIM-40 format, the VPT-40 targets industrial applications like machine vision, medical imaging, and process automation. The module can be used with any carrier card that provides TIM-40 sites.

The VPT-40 can acquire and display images and add overlay data. Featuring horizontal and vertical blanking, the module accepts one of four software-selectable NTSC, PAL, RS-170, and CCIR monochrome, interlaced video inputs. The input resolution can be up to 640 x 480 (NTSC) or 690 x 580 (PAL) at 8 bits per pixel. The VPT-40 produces an RGB (RS-343A) video output. The output resolution can be 640 x 480 or 690 x 580 at 24 bits per pixel with a 4-bit,15-color RGB overlay.

The video-input mode, video format, image size, RGB overlay, digitizer-lookup RAM data, RGB color-palette RAM data, digitizer gain, and offset are all software selectable. To provide high-speed access to video data, the module provides three 8-bit wide TMS320C4x-compatible comm ports. The VPT-40 also has seven 75-R SMB-type connector jacks (4 input, 3 output).

The VPT-40 costs $1995 in single quantities.

Ariel Corp.
2540 Rte. 130
Cranbury, NJ 08512
(609) 860-2900
Fax: (609) 860-l 155
http://www.ariel.com/

#501

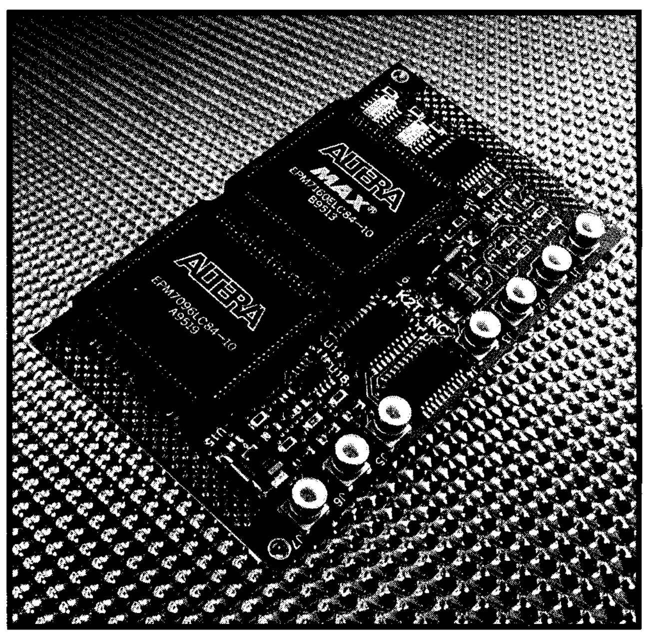

NEW PRODUCT NEWS

PENTIUM SINGLE-BOARD COMPUTER

Advantech announces a Pentium-based PCI/ISA single-board computer. The PCA-6157 SBC features a choice of 75–150-MHz Pentium CPU. It is available with a PCI local bus or ISA bus, 256-MB DRAM capacity, 256-KB external cache expandable to 512 KB, a PCI SCSI-II controller, four PCI enhanced IDE HDD interfaces, two FDD interfaces, two RS-232 serial ports, and an ECP/EPP-compatible parallel port. The PCA-6157 also features an onboard keyboard connector, real-time clock/calendar, 12-level watchdog timer, and power-on self-test diagnostic LEDs.

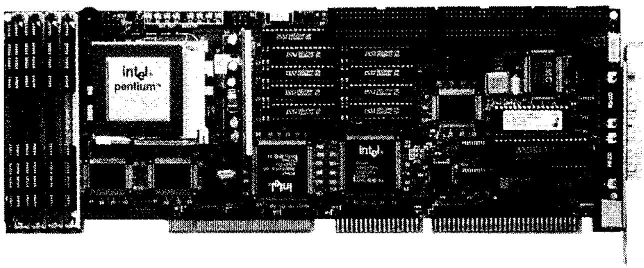

The Pentium SBC meets the tough challenges of an industrial environment. It's ideal for heavy-duty processing tasks such as image or voice processing, high-speed serial communication, medical or pharmaceutical applications, and network server.

The PCA-6157 price range starts at $1050 with a 75-MHz Pentium CPU.

Advantech
750 East Arques Ave.
Sunnyvale, CA 94086
(408) 2456678
Fax: (408) 245-8268

#502

GRAPHIC DISPLAY MODEM

Model 227 Speed Link from Telebyte Technology provides high-speed, full-duplex transmission from DC to 115.2 kbps over two twisted pairs. It incorporates the DataSpy LCD display, which provides on-line status of Transmit Data, Receive Data, and five control signals. The display uses less than 1 mW of power and doesn't compromise Model 227's operation.

The addition of the LCD display means that the user can view the instantaneous performance of the data communications link. This capability simplifies and speeds installation and functional checkout of Model 227. It also assists in troubleshooting any future networking problems.

Model 227's speed capability provides the basis for a communications link for attaching remotely located PCs to a LAN. This operation is accomplished by installing a terminal server on the LAN and using Model 227s to extend the high-speed serial data lines to the remote PCs.

Model 227 does not require a power source. It steals its operating power from the Transmit Data signal applied to its RS-232 connector. When the LCD becomes active, the user knows that sufficient power has been applied to the modem. Pins 2 and 3 of the RS-232 connector provide data I/O. The DTE/DCE switch reverses pins 2 and 3, depending on whether the modem is connected to a DTE or DCE device.

Model 227 is offered with a male or female DB-25 RS-232 connector. The four-wire line interface is available on a five-position terminal block and an RJ-11 modular phone connector. This four-wire line interface includes built-in surge protection which uses pin 1 (Frame Ground) as the path for suppressing surges.

Model 227 is packaged in a rugged plastic case measuring 2" x 4.25" x 0.75". It sells for $109 in singles.

Telebyte Technology, Inc.
270 Pulaski Rd. • Greenlawn, NY 11740
(516) 423-3232 • Fax: (516) 385-8184
sales@telebyteusa.com

#503

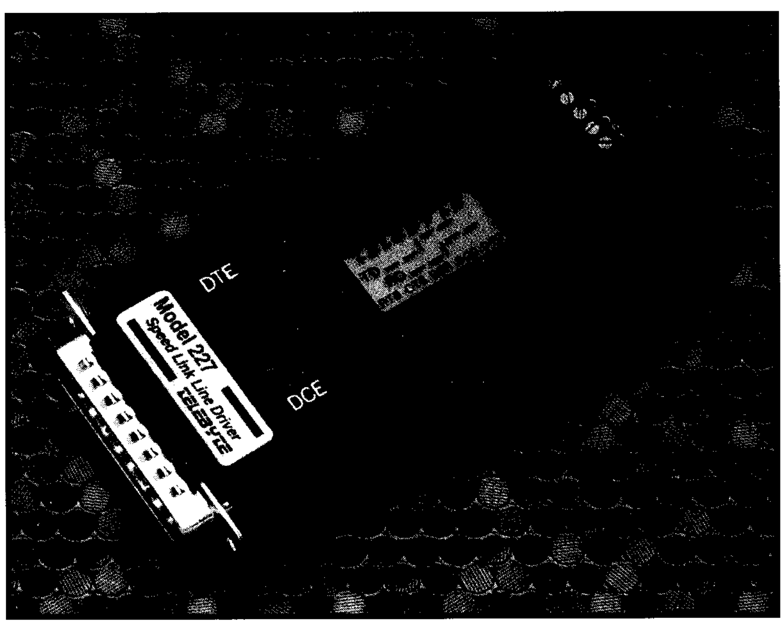

NEW PRODUCT NEWS

APPLICATION BUILDER SOFTWARE

DataViews has announced **DV-Xpresso,** a fully featured application builder from which custom graphics can be built to monitor and control complex processes in dynamic applications. DV-Xpresso integrates a GUI builder, a 3D Graphics Builder, suites of 2D or 3D graphs, control widgets, special graphs, and a table widget. This combination enables programmers to build custom graphics for X Windows in an object-oriented architecture without hundreds of tedious function calls.

There are over 50 meters, dials, and knobs to choose from in building custom control panels. These controls are highly interactive and support several forms of user feedback. They react to a mouse click whereby a displayed value is changed and an active element (such as a needle or an arrow) is moved.

With the complete 3D graphics editor, programmers can build a customized application with 2D or 3D graphical objects representing an animated graphical drawing used in an application. From this editor, parts of the widgets or the widget itself can be saved to create libraries of reusable graphical components.

The DV-Xpresso graphics builder lets you build animated graphical drawings which, when integrated into a widget, provide a complete interactive environment for process control, monitoring, and simulation applications. Such applications include chemical, manufacturing, network monitoring, air and traffic control, and process simulation. Some of the features available include real-time 3D shading and hidden-surface removal (using the Z-sorting technique), grouping and traversing object hierarchy, color and font tables editable on a per-widget basis, double buffering, and many others.

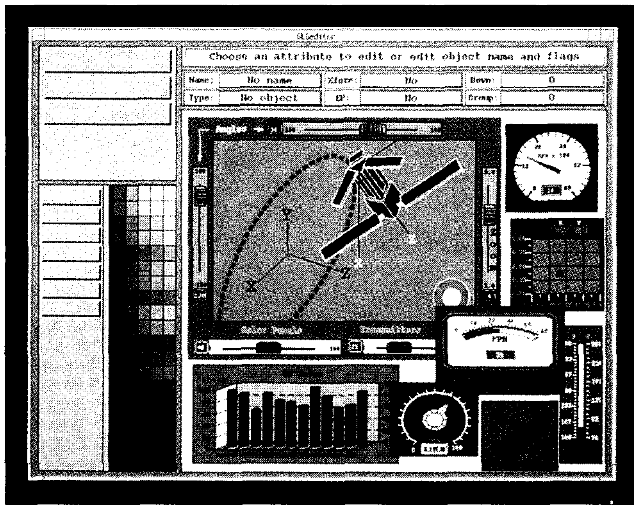

DV-Xpresso is available on the SunOS, Solaris, HP/9000, Silicon Graphics IRIX, QNX, and Windows NT. Alpha Digital UNIX, RS/6000, and Windows versions will follow. Prices start at **$1000.**

DataViews Corp.
47 Pleasant St.
Northampton, MA 01060
(413) 586-4144 • Fax: (413) 586-3805
infoQdvcorp.com

#504

COMPUTER MONITOR CHECKER

Data Sync has introduced a series of compact, low-cost computer monitor testers that are used for verifying the operation of fixed or multiscan monitors. Models are available for MDA, CGA, EGA, VGA, SVGA, and XGA monitor types and are offered with either an outlined color-bar pattern or a sequenced red-green-blue screen.

The displayed image is made up of 15 color bars outlined with a white border. The pattern size simulates the display area of most programs. Symptoms can usually be identified by the displayed pattern.

The RGB tester models continuously sequence through red, green, and blue screens and are useful for checking color purity and performing extensive burn-in testing. Each color is displayed for approximately four seconds.

Each model is equipped with a DC power jack and the appropriate female D-sub connector. The tester measures 2.25" x 2.18" **x 0.65"** and is powered from a supplied 9-VDC wall adapter. Prices start at $24.95.

Data Sync Engineering
40 Trinity St.
Newton, NJ 07860
(201) 383-I 355
Fax: (201) 383-9382

#505

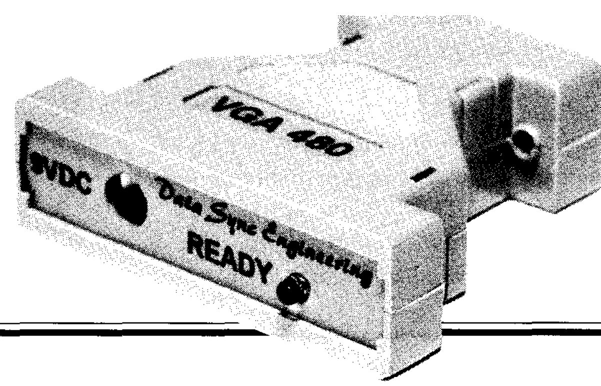

NEW PRODUCT NEWS

CAPACITOR DATA BOOK

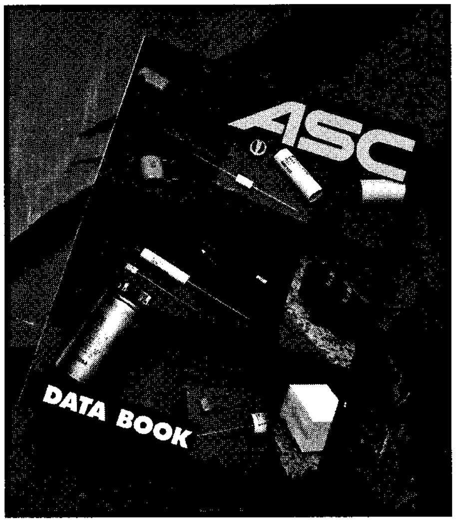

ASC Capacitors has released a new **data book** summarizing its product offerings and providing useful information for anyone interested in capacitors. The data book provides general information about specifying the correct capacitor type and value. Topics range from the basic "What are capacitors?" to more advanced topics such as "How close is a close-tolerance capacitor?" "Capacitance vs. temperature," "Insulation resistance and leakage current," and "Dielectric absorption." For further reading, the book lists appropriate technical papers, also available from ASC.

The data book incorporates a section on application-specific or custom solutions. The section discusses which capacitor parameters can be altered to produce solutions that are unavailable in standard configurations. Among the many solutions mentioned are extremely high-voltage or very low E.S.R. capacitors, custom AC- and DC-based film formulations to enhance dielectric performance, and completed capacitor subassemblies to reduce manufacturing cycle time and inventory costs.

The data book also catalogs the company's more than 40 standard capacitor lines. The publication enumerates physical and electrical characteristics, and includes drawings to clearly illustrate specifications. Types of capacitors described include those using polypropylene, polycarbonate, polyester, polystyrene dielectrics, and polyester RC snubber networks.

ASC Capacitors
301 w. **"O" St.**
Ogallala, NE 69153
(308) 284-3611 • Fax: (308) 284-8324 **#506**

#104

#105

NEW PRODUCT NEWS

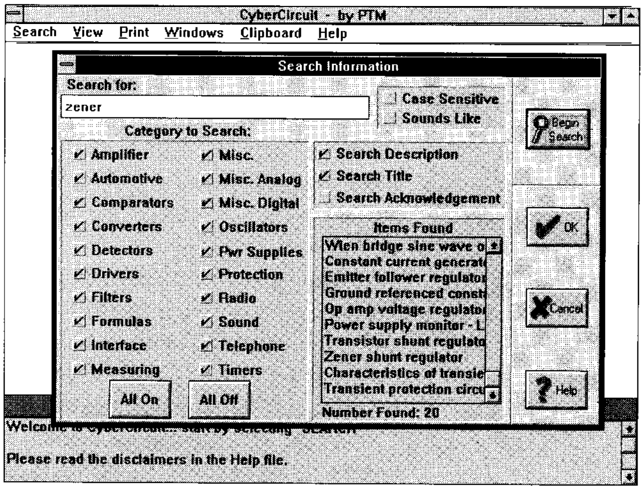

ELECTRONIC CIRCUIT DATABASE

CyberCircuit from PTM is a Windows program that contains hundreds of circuits along with formulas, IC dimensions, interface pinouts, dB tables, DTMF frequencies, thermal information, graphs, and charts. The circuits include power supplies, amplifiers, filters, timers, oscillators, drivers, and more.

The program eliminates the need to hunt through reference books or application manuals. The circuits were gathered from electronic industry data books, application notes, and publications. Proprietary information and circuits from PTM are also included.

CyberCircuit features a fast, powerful search engine. Searches can be done on titles, acknowledgments, or complete text. Circuits can be searched by part number, function, source of information, or any other key word. In addition to the desired schematic, circuit descriptions, part values, and references can be retrieved as well. A "sounds like" feature compensates for unknown or incorrect spelling. An extensive help system is included. Graphics are redrawn rather than scanned. Search information can be printed on any Windows-installed printer.

PTM
119 Benton Oaks Dr.
Sauk Rapids, MN 56379
(612) 253-0591
Fax: (612) 253-0591

#507

NEW PRODUCT NEWS

I/O SLAVE BOARD

The **MPC 486** is a complete single-card '486 that interfaces to a host motherboard through a 128-KB block of dual-ported, bidirectional shared memory. This design offers high-speed interprocessor transfer of large amounts of data.

Multiple boards can be used with one host—an arrangement similar to having multiple workstations and a LAN inside one PC. Each board can run a different application or even operating system.

The MPC 486's PC-compatible architecture means all popular languages and development tools can be used. Any tool that generates a PC-executable file (.COM or .EXE) can build an MPC application. By connecting a monitor and keyboard to the MPC 486, debuggers can be run native.

The MPC 486 includes onboard 12-bit DACs, high-power parallel I/O, and external interrupt support. The board also provides COM1 and COM2 serial ports, LPT printer port, real-time clock, floppy and IDE controllers, and keyboard support.

Onboard PC/104 and iSBX expansion connectors support the broad selection of add-ons, including additional I/O, LANs, avionics, DSP, CRT and LCD displays, speech, multimedia, GPS, and motor control that are available from over 100 manufacturers.

The MPC 486 I/O slave board sells for $1495.

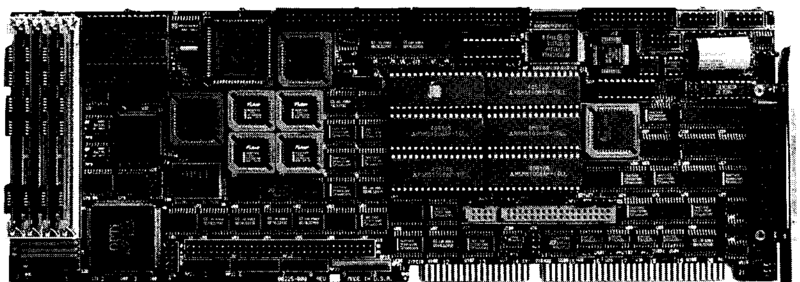

Technological Artisans, Inc.
145 West 67th St., Ste. 6E
New York, NY 10023
(212) 721-9769

#508

FEATURES

FEATURE ARTICLE

Brian Millier

Capacitive Measurement of Fluid Level

Need to monitor a remote tank? If so, listen up. Brian modifies an idea from a ConnecTime discussion to come up with a microcontroller-based readout with control capability and a trend feature.

I wonder if many *INK* readers are "backwards" like I am. Sure, I look at the "Inside Issue" page first, but I then often turn to "Priority Interrupt" and the various threads in "ConnecTime."

I read this way because I pick up my mail en route to work and have a bit of time to read during the commute. Short topics are appropriate. More importantly, though, I often find that by keeping snippets of information I've read in the back of my mind, I'm able to apply them to problems or opportunities at work that come up during the day.

A case in point concerned a thread in *INK 45* about the measurement of water level. Two techniques discussed were a refraction-of-light method and a capacitive method (contributed by Pellervo Kaskinen). At the time, I needed to replace a failing water-interlock sensor on a large still.

This conductivity-type sensor was prone to fouling and sometimes failed in a mode which caused the heating elements to burn out. I fabricated a small light-refraction tube as described in that thread. See the sidebar on "Refractive-Index Sensing" for details.

However, the same thread contained another reference. It discussed the idea of measuring water level by monitoring the change in capacitance if the liquid itself were used as the dielectric material of a capacitor.

I remembered this concept recently when we needed to remotely monitor the distilled-water level in the 200-gal. storage tank at my work place. On numerous occasions, large amounts of distilled water were wasted when valves were left open.

The tank was already fitted with a sight glass (made of ⅝″ thin-wall clear plastic) running the entire 36” height of the tank. This technique gave a visual indication of water level. However, the tank was located in a remote mechanical room in the building's attic.

I needed to turn the sight glass into a capacitor and design a simple microcontroller-based readout with control capability and a trend feature. Let me tell you how I did it.

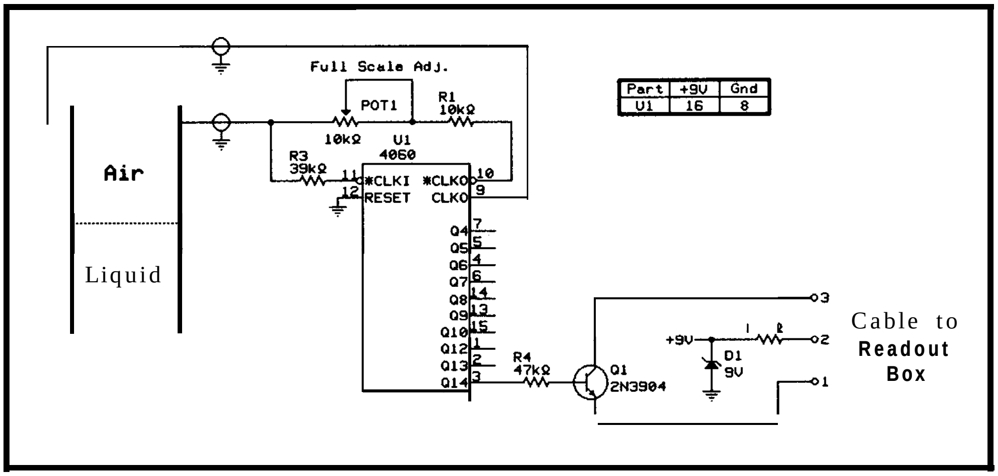

Figure 1--Due to the very low capacitance of the sensor, the oscillator head is built separately and placed directly at the sensing site.

THE CAPACITOR EXPERIMENT

We all use capacitors in circuitry and have a reasonable understanding of how they work. The actual details of capacitor construction and the dielectric constant properties of different materials are probably topics you haven't spent much time thinking about. Certainly, I haven't.

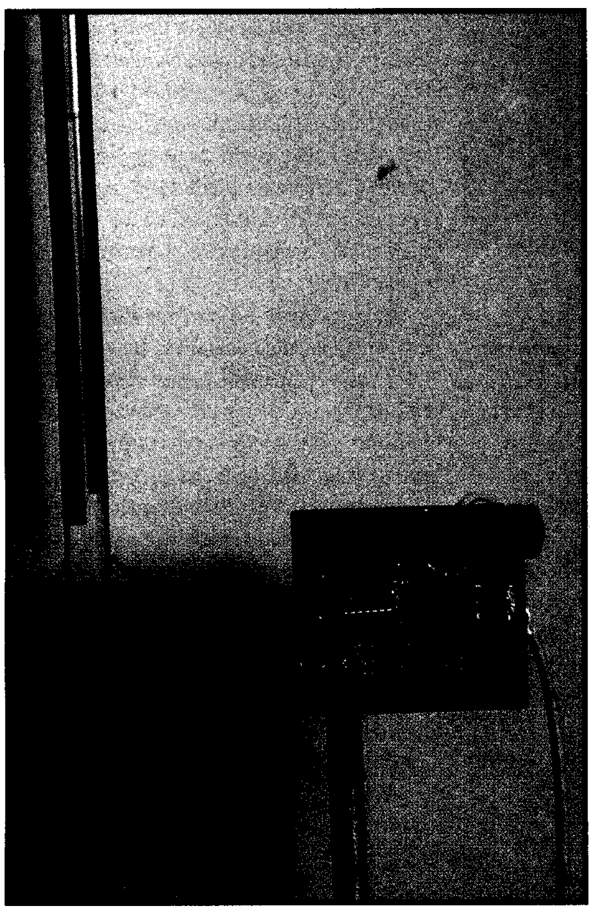

Photo 1—A closeup of the oscillator head circuit shows the short coaxial cables connecting it to the capacitive sensor.

Intuitively, I felt that any capacitor using a sight tube of the previously mentioned dimensions would have a value somewhere in the picofarad range. And, I expected that the change in capacitance occurring as the dielectric went from air to water wouldn't be too great. I also worried that the Q (quality factor) of the capacitor might be low when using water as a dielectric.

I first fabricated an experimental capacitor. It consisted of a ½″-diameter, glass graduated cylinder, 50 cm long. If you took chemistry, you undoubtedly used one. As you may recall, the volume of the liquid contained in them can be read from graduations along their length.

The capacitor plates were made of ¼″ copper foil with an adhesive backing. This material is available in stores catering to stained-glass craft. I attached two strips of this foil the full length of the cylinder. Each strip was centered 180” away from the other on the cross section of the cylinder.

When I connected this capacitor to a capacitance bridge, I noted:

- the capacitance was in the 20–50-pF range, depending on the level of the water in the cylinder
- the Q of the capacitor was low, resulting in a very broad null on the capacitance bridge. Practically speaking, it was impossible to get a reading with any amount of resolution.

From these observations, I made some conclusions about possible measurement techniques. Due to the low capacitance, it's essential to position the measurement circuit as close as possible to the sight-glass capacitor to minimize stray capacitance.

Also, the low Q affected the type of circuit I could use to measure capacitance. Since the tank was in a remote location, the method I used to measure the water level in the tank needed to be able to accurately transmit over several hundred feet via inexpensive cable such as telephone wire.

Figure 1 shows a schematic diagram of the oscillator circuit I settled on. It was built first and then connected to a frequency counter early in the design phase. Photo 1 gives a closeup view of the circuit.

The heart of the circuit is a CMOS 4060 IC oscillator and 14-stage divider IC wired as an RC circuit. The sight-glass capacitor (SGC) forms the C component. Rl and pot R2 form the R component.

Due to the low capacitance of the SGC, oscillation frequency is in the order of several hundred kilohertz. I found that the 4060 oscillated quite well in this configuration, even given the low Q of the SGC. The relatively high frequencies led to these design considerations:

- the 4060 operates on 9 V rather than on the standard 5 V. CMOS ICs work better [if they work at all) at high frequencies if the power supply voltage is increased toward the +15-V limit.
- short coaxial cables (RG- 174 type) interconnect each SGC plate to the 4060 oscillator
- the coaxial cable connecting the SGC to the 4060 pin 9 solders directly to the bent out pin of the 4060 [rather than connecting through the socket] to minimize stray capacitance

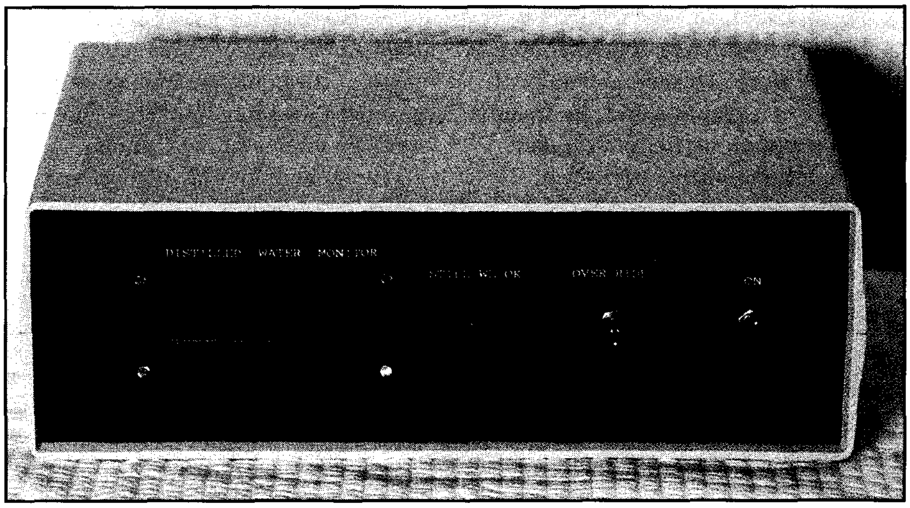

Photo 2—*The* display unit can be conveniently located where it is easy to monitor. It connects to the oscillator head *via a 3-wire cable.*

To supply a low-impedance TTL-compatible drive signal, I used Ql as an open-collector driver with its load resistor returned to 5 V on the readout board.

A 9-V zener diode regulates the raw DC voltage provided by the readout module since the 4060 oscillation frequency varies somewhat with power-supply voltage.

Figure 2 is a graph of the oscillator period (multiplied by the chosen divider value of 2^{14}) versus water level in the SGC. It turns out to be reasonably linear, but there is a significant oscillation period even when no water is present in the SGC.

This oscillation occurs because the air-filled sight tube itself has a significant capacitance. The stray capacitance of the connecting wires also has an effect.

From the experiment results, I decided that the concept had merit. I chose the CMOS 4060 for its built-in divider chain and its ability to act as a variable-frequency oscillator using the capacitance available from the SGC.

To transmit the water-level signal over a long distance, I felt that a low-frequency pulse train would be less prone to transmission errors than other methods.

The multistage divider chain in the 4060 was ideal. It took whatever oscillation frequency the SGC produced and, by choosing the correct divider tap, gave a low-frequency pulse train in a range of less than 100 Hz.

This range can now be measured by microcontrollers such as Motorola's 68HC705K1 or Microchip's PIC.

HARDWARE DETAILS

Figure 3 gives the readout and control circuitry, which can be located anywhere the user finds convenient. The SGC oscillator, which must be located at the tank itself, connects to the readout via a three-conductor cable.

If you only had two wires [a preexisting cable run, for example), the circuit could easily be changed to accommodate that. Then, Ql (the oscillator output buffer) could modulate the current drawn by the oscillator. This current difference could provide the signal at the readout end.

The circuit is quite simple. However, it includes one interesting feature. Many circuits require some form of user-adjustable calibration control(s).

Generally, this adjustment is performed once or very infrequently and must survive after power is cycled on and off. In analog circuits, it is best done with a potentiometer-possibly a ten-turn type if resolution and stability are important.

If the signal being measured is digital, it doesn't lend itself to an analog solution such as a pot. In this case, the calibration can be done using actual numbers stored in the microcontroller's memory. This method works well if the circuit contains a microcontroller with some form of nonvolatile memory (such as EEPROM or battery-backed RAM), a keypad or switches for input, and a display to prompt the operator.

As you cram controller functions into increasingly smaller, less expensive microcontrollers, the shortage of EEPROM and I/O port lines for switches make calibration of parameter input and storage difficult. In this design, I needed to calibrate a value for the oscillation period with the tank empty. In another circuit I designed recently, a calibration value was needed to convert the digital value obtained by measuring wheel revolutions to distance traveled, given various wheel diameters.

The solution I used is incorporated in the circuit shown in Figure 3. The calibration pulse generator consists of a 555 timer IC configured as a monostable multivibrator. The width of the pulse it produces at output pin 3 is proportional to the RC time constant

Refractive-Index Sensing

The refractive-index method of sensing liquid level provides a low-maintenance sensing alternative. It doesn't rely on the liquid's conduction of electricity, which is susceptible to sensor fouling after time unless the liquid is pure. It also doesn't need a float switch or any other mechanical device which might seize up or rust.

The device consists of a clear cylindrical rod with a rounded end. I used 3⁄16″ solid glass tubing, since glass tends not to get fouled by impurities in the water. I understand that a similar cylinder of clear plastic with the end cut at a 45” bevel also works.

The principle at play involves the refraction of light at a boundary between two dissimilar materials. Light is sent down the cylinder. If the refractive index of the material on the other side of the rounded end is much different from the refraction index of the tube itself, the light is reflected back. This is true when the rounded end is exposed to the air. However, if you immerse the end in water, which has a similar index as glass, light passes through, reflecting very little.

To send out this light and monitor its reflection, I used an inexpensive Optek OPB706A, a common reflective object sensor assembly which contains an IR LED and a phototransistor sensor. The unit is a ¼″ cube that matches up well with the small glass cylinder. It's important to couple the light-emitting-and-detecting face of this device directly to the end of the glass cylinder.

I fashioned a small enclosure using a piece of ⅝″ aluminum cylinder stock. Once the cylinder and OPB706A are properly mounted, silicone cement holds them in place.

I kept the length of the glass cylinder to less than 1 .5″ to minimize loss of light. The aluminum enclosure also keeps any stray light away from the detector. The rounded end of the cylinder must be in contact with the liquid. So, insert it into a hole in the tank and seal it, if that's the type of measurement you're doing.

When I connect the phototransistor collector to the circuit as shown in Figure i, the voltage varies from about 7 V in water to 3 V when there's no water left. Therefore, I set the threshold potentiometer to 5 V, which gives a stable output reading from the LM339 comparator as the water level changes.

With a different light cylinder or opto device, you need to determine the optimum setting of the threshold pot. Similarly, if an output signal of the opposite sense is more appropriate for a given design, you can easily reverse the connections to the positive and negative pins of the comparator. I made use of an available 12-V supply, but the whole circuit could be modified to work on 5 V.

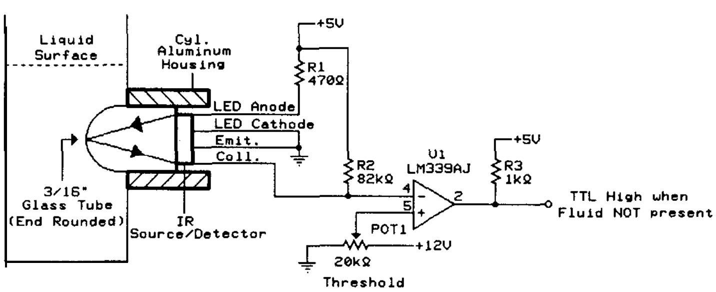

Power Table

Part	+12V	Gnd
U1	3	12

Figure i—*An* **optical** *sensor and comparator make up a simple refractive-index fluid level sensor.*

of C2 and the series resistance of R6 and pot R4.

The microcontroller uses the PB1 I/O pin to monitor the 555 pulse and measures the width using a software loop. The 555 is triggered using the PBO I/O pin.

Typically, that pin can be used for other purposes as it is only needed for calibration once at powerup. In this design, PBO provides the drive signal to the control relay at all other times.

In operation, the microcontroller triggers the 555 early in the program sequence and measures the pulse width. It displays this value on the LCD display momentarily (for user feedback), and then the program proceeds.

The user would normally adjust R4, reset the microcontroller, and observe the number until the desired calibration value is reached. This circuit provides a nonvolatile, easily-adjusted calibration block for about $2 using only one dedicated I/O pin. I use a nonelectrolytic type of capacitor for C2 to ensure stability.

The rest of the circuit operates as follows. The low-frequency pulse train arrives at pin 3 of J1. R3 acts as the collector load resistor for Q1 on the oscillator board. R6, D2, D3, and C6 are included for protection and signal conditioning.

A 7414 Schmitt trigger ensures a nicely squared-up signal and provides noise immunity. This 7414 would not be necessary if the cable run to the oscillator box were short or shielded from noise.

The clean pulse-train signal is fed into the *IRQ pin of the microcontroller. The period of that pulse train is measured as described in the next section.

This circuit contains a control relay (K1) with an associated driver (Ql) fed by PBO. The software has provisions for actuating this relay when the level drops below a given value, and deactivating it when the level in the tank approaches full. This circuit shares PBO with the calibration circuitry, as the two functions are mutually exclusive.

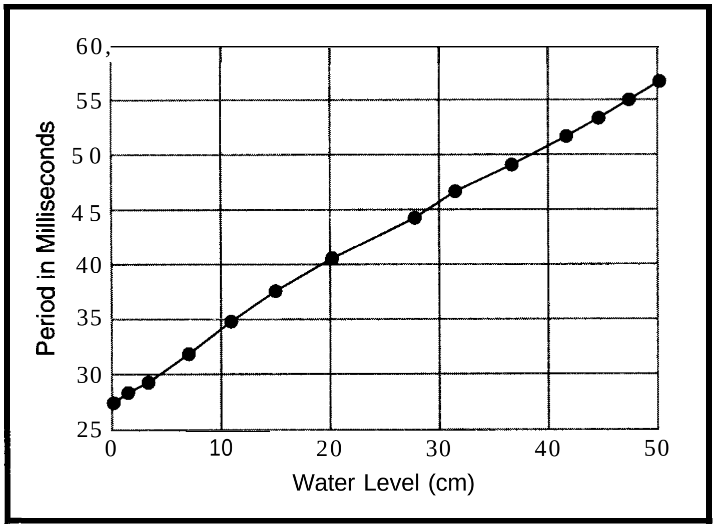

Figure 2—*The oscillator/divider* output period **varies *with* the water level obtained using a 50-ml glass graduated cylinder as the sight-glass capacitor.**

The LCD display is a common 16 x 2 display available for $8 from suppliers such as Timeline. It runs in the 4-bit interface mode. The driver software is described in more detail in my article "An LCD and Keypad Module for the SPI" in *INK* 57.

A simple full-wave rectifier and 7805 provide the V_{CC} for the readout module. They also provide a raw +10 VDC for the oscillator circuit.

My application called for control of the tank liquid level, which depended on the availability of AC power. However, if your application just involves monitoring a liquid level, the relay circuitry can be deleted. The whole circuit would then draw less than 10 mA, allowing it to run on a battery supply if so desired.

FIRMWARE

The firmware for this project is quite simple since only 496 bytes of EPROM are available in the 68HC-705K1 microcontroller. In this design, I omitted my customary LCD signon message, which identifies the unit and shows my company logo, as there was no room left in the EPROM to accommodate it! The complete, commented source code for the firmware is available on the Circuit Cellar BBS.

The calibration function, outlined in "Hardware," consists merely of triggering the 555 monostable and entering a software timing loop which polls PB1 to detect the end of the pulse.

I actually trigger the 555 twice. A dummy trigger, which is not used, is followed by another trigger. Only the pulse from the second trigger is measured because the first time a 555 is triggered, its output pulse is different (and less accurate) than subsequent triggers.

The SGC oscillator period is measured by counting the total number of timer overflows of the 68HC705K1 internal timer that have occurred from one *IRQ negative edge to the next. The timer overflow rate (with the 3.2768.MHz crystal used) is 0.625 ms.

Due to the limited program memory space and math instructions in the 68HC705K1, the algorithm used to convert the pulse period to liquid level had to be carefully chosen. Therefore, I broke down the range (0–100%) of the liquid-level display into 2% increments, using the following formula:

liquidlevel(%) = 2 x (#timeroverflows – offset)

where offset is the zero value obtained from the calibration circuit on power-up.

It follows from this formula that there must be 50 more timer overflows for the full tank (100%) reading'than for the empty (0%). Since timer overflows occur at a 0.625-ms rate, the difference in oscillator period needed is 50 x 0.625 ms or 3 1.25 ms. I describe how to achieve these values from the SGC oscillator in the next section.

The offset parameter is needed because the oscillator has a significant period of oscillation even with the sight glass (and tank) empty. It must be subtracted to make the display read zero under this condition.

I implemented a real-time clock using the real-time interrupt capability of the 68HC705K1's timer block. This clock provides an hour:minute:second readout on the LCD updated each second. While a clock display is not absolutely necessary, I included it because I also needed a way of measuring the elapsed time to support the trend function.

The trend function provides for the display of liquid level one and two hours ago, as well as the current reading. It provides the operator with an insight into trends and problems such as leaks. Every minute, the liquid level is compared to two setpoints.

In my application, I wanted to turn on a large still when the water dropped below 35% and turn it off when it reached 70%. (Another small, noncontrolled still topped up the tank beyond 70% .) So, I drive relay R1 using the PBO output line. Relay R1 contacts brought out to connector J2 drive a large contactor in the still.

Once the program completes initialization of variables, LCD, and performs the calibration pot reading, it enters a low-power wait state. All operations after that are initiated by interrupts-the 'IRQ interrupt of the oscillator pulse train or the timer interrupts associated with the real-time interrupt and timer overflow. Most of the 32 bytes of RAM contained in the controller are used by the program variables and stack.

CONSTRUCTION

To construct the measurement device, you first need to make up the sight-glass capacitor. If a sight glass is not already installed on the tank, add one. The mechanical details are left to you, but I'd recommend using thin-walled glass or plastic tubing. I got satisfactory results with both glass and plastic tubing 0.5–1.0″ in diameter and I-3' in length.

For tubing diameter up to 0.625″, a single width of ¼″ copper foil works well for each capacitor plate. Beyond that, a wider plate is better. Apply the foil on opposite sides of the sight tube, being careful to apply them evenly.

The next step is to construct the oscillator and connect it to the sight-glass capacitor. Photo 1 gives a close-up view of the oscillator head connected to the sight-glass capacitor on our 200-gal. tank.

It's easiest to connect a frequency counter (or scope) and then measure the oscillation period. Start by connecting to the Q14 tap of the divider. You'll find that the oscillation period with the sight tube empty is about 20-50% of the period when it is filled with water.

The difference between these two readings must be adjusted to 31.25 ms. If your readings are fairly close to this, you merely adjust R2 on the oscillator board. If the difference in period readings is considerably greater than 3 1.25 ms, pick another tap on the divider chain to get into the right ballpark and then adjust R2.

Keep track of the empty and full periods after you complete the adjustment. Since the oscillation is at a fairly high frequency and the sight-glass capacitor Q is low, monitor the divider output on a scope while allowing the liquid level to go from full to empty.

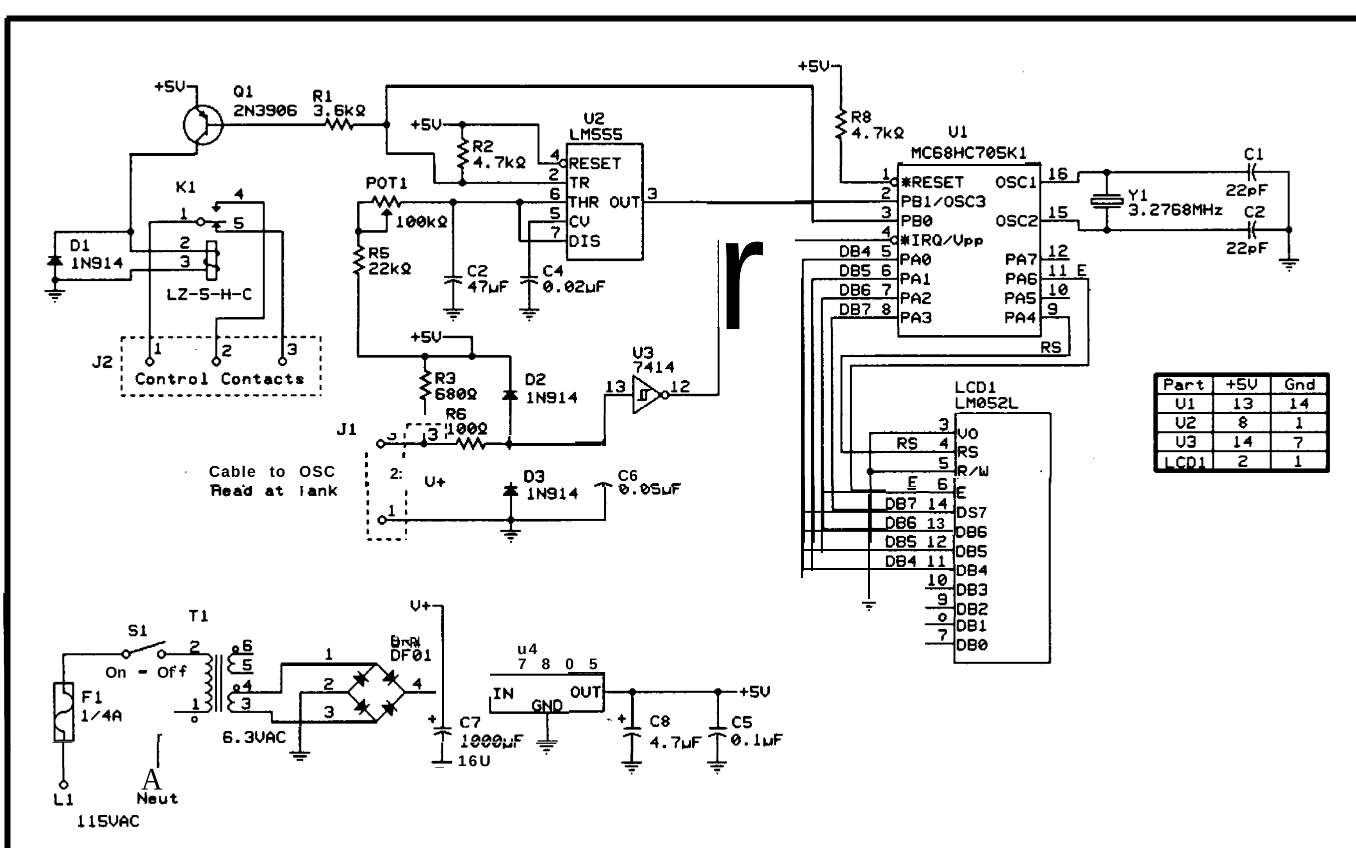

Figure 3—*The* **wafer-level monitor is** *based on an MC68HC705K1* **processor** *and includes a small LCD display. The circuit is located in a cabinet at the* **monitoring site.**

This procedure ensures that the oscillator is working properly throughout the range. I can't guarantee that all potential physical layouts you want to use will work.

Assuming that the empty-tank period is somewhere in the 20–50% range as compared to the full-tank period, the zero pot adjustment on the readout board corrects for this zero offset.

By completing the above steps, you verify that the concept works in your particular application. I haven't tried this approach on very tall tanks or with liquids other than water.

After the readout module is built, connect it to the oscillator board, and power them both up. For one second, a single number appears on the LCD display for diagnostic purposes.

This number varies depending upon the setting of pot R4. It is used as the offset value mentioned in formula 1 in the firmware section. Basically, R4 must be adjusted until this number is equal to the tank empty period (in ms) divided by the timer-overflow rate (0.625 ms).

Practically speaking, it is easiest to adjust pot R4 to obtain an LCD level reading of zero with the tank empty. Therefore, make sure the tank is empty and then cycle the circuit power on and off repeatedly, adjusting the R4 pot a little bit each time until a zero reading is obtained.

If you've adjusted the oscillator as specified, the LCD should now accurately indicate liquid level without further adjustment. If not, go back and tweak R2 on the oscillator and R4 on the readout.

You'll see that the LCD displays three level readings on line one: Current, Last Hour, and Two Hours Ago. The latter two readings are zero until one and two hours have passed. Line two of the LCD displays the time, updated every second. Photo 2 presents the display unit from my system.

WRAP UP

A few drops of water can be your best friend when you're out in the desert, parched and starving. However, an overabundance can be your worst enemy when it floods everything.

I hope you can take these ideas and techniques and use them in your own applications. ■

I'd like to thank Pellervo Kaskinen for the concept he described in the "ConnecTime" thread. I also want to thank Danny Abriel of Dalhousie University for the photos.

Brian Millier has worked as an instrumentation engineer for the last 12 years in the chemistry department of Dalhousie University, Halifax, NS, Canada. *In his* leisure time, *he operates Computer Interface Consultants and has a full electronic music studio in his basement. He may be reached* at *brian.millier@dal.ca*

SOURCES

MC68HC705K1P
Jameco Ltd.
1355 Shoreway Rd.
Belmont, CA 94002-4100
(800) 83 l-4242 USA
(415) 592-8097

LCD module
Timeline, Inc.
23650 Telo Ave.
Torrence, CA 90505
(310) 784-5488

MC68HC705K1/D
68HC705K1 Technical Data Manual
Motorola Literature Distribution
P.O. Box 20912
Phoenix, AZ 85036
(602) 244-6900

Programmed **68HC705K1P** ($15)
Computer Interface Consultants
P.O. Box 65 Site 17
R.R. #3 Armdale, NS
Canada B3L 4J3
(902) 876-8645
brian.millier@dal.ca

I R S

401 Very Useful
402 Moderately Useful
403 Not Useful

Lightning Protection Techniques

FEATURE ARTICLE

Leon Byerley

Traditional lightning protection fails to adequately protect the complex electrical infrastructure we now have. Leon introduces us to the preemptive protection currently used to improve electrical systems' reliability.

Each year, lightning causes about a billion dollars' worth of damage to residential insurance customers in the United States [1]. Annual lightning-caused losses to industry and commerce are not known but could total billions of dollars with damages and downtime.

Lightning causes the majority of commercial power sags, interruptions, and outages [2] which interrupt and upset continuous or critical industrial processes. Clearly, the practice of lightning protection (including surge suppression, EM1 control, and electrical protection in general) has not kept up with technology.

This fact is evident from a quick look at ANSI/NFPA 780, *The Lightning Protection Code*. It provides minimal instructions for comprehensive shielding, grounding, bonding, and the coordination of surge-suppression devices in a lightning-protection scheme.

Lightning protection is usually an afterthought or retrofit and the results are often a mixed bag. With the exceptions of telecommunications and electric-power utilities, few companies do thorough lightning-protection engineering "up front."

Fortunately, modern airplanes are well engineered to withstand direct lightning. For a model of how to do lightning-protection engineering for ground-based equipment, look no further than a modern airliner.

TRADITIONAL PROTECTION

Lightning protection, conceived and proven by Benjamin Franklin around 1750, aimed to protect structures and their occupants. The efficacy of Franklin's prescriptions for protecting structures from lightning was immediately noted.

Franklin discovered that lightning is an electrical phenomenon and that structures can be protected from the effects of lightning. As a result of his discoveries, he was elevated to a position of worldwide prominence. However, in Franklin's time, there was no electrical infrastructure to protect from lightning.

Modern lightning-protection codes reveal that lightning protection has advanced little in 250 years. A deeper look shows that common lightning-protection practices are not a scientific or engineering discipline, but rather an art handed down from generation to generation.

Lightning protection, as described in codes such as ANSI/NFPA 780, UL96A, and LPI 175, is not based on physics and first principles but on historical precedent and empirical evidence.

There are no papers in the peer-reviewed scientific or engineering literature that show how or how well traditional lightning-protection technology performs under the stress of natural lightning. It's no wonder that in the lightning-protection field, snake-oil salesmen have been at work for 250 years.

Scientists still don't understand the behavior and efficiency of the traditional sharp-pointed lightning rod designed by Franklin. Nor do we understand the surge impedance of down conductors and ground electrodes when energized by lightning.

Since the surge impedance of the ground electrode is not understood, we cannot understand the performance of surge-suppression devices, which are entirely dependent on this variable for proper function.

This lack of knowledge doesn't mean that traditional lightning protection does not work or that it's wrong to use UPS and surge-suppression devices. But, it does mean that there is no quantitative understanding of how well lightning-protection systems function. Understandably, testing with natural lightning is not easy.

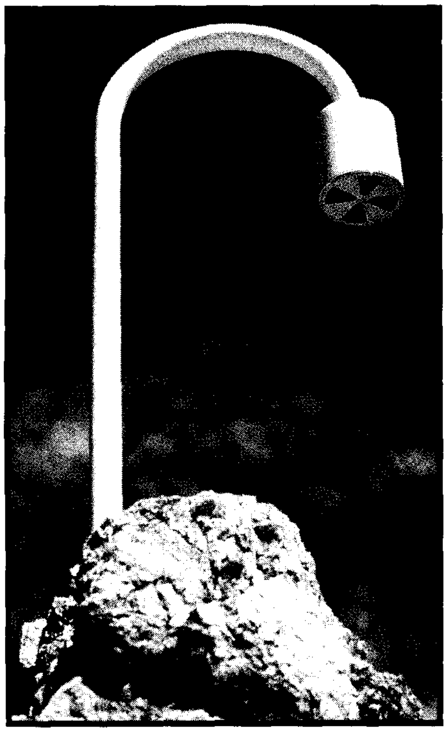

Photo 1-A field *mill* measures *the* static and *slowly* changing components of electric fields associated with a thunderstorm.

We know that traditional air terminals (lightning rods) and ground electrodes don't always work optimally. While we understand the electrical properties of the Earth at steady state, we know almost nothing about what happens in and on the Earth during the extremely energetic and dynamic situation of a lightning discharge to ground. See the sidebar "A Simple Model" for an example of how standard ground rods work when struck by lightning.

We've recently learned that when lightning current is injected into an array of long vertical ground rods in soil of high conductivity (a "good ground"), most of the lightning current appears on the Earth's surface in the form of surface arcs and fireballs over 20 m long [3]! For lightning discharges, the resistance (*R*) of the earth connection appears to make only a small contribution to the impedance of a ground electrode.

Advertising informs us that liberal application of UPS and surge suppression devices solve lightning problems. As a result, most of us can't decide rationally what electrical protection measures are appropriate for computing machinery, business machines, and home-entertainment equipment in a middle-class household.

PREEMPTIVE PROTECTION

Some recent lightning-protection schemes involve more than the traditional protection components (surge suppressors, lightning rods, ground electrodes, and so on). They also include some type of automatic or manual preventive action when lightning threats occur. Preemptive protection is used in a variety of settings, in some cases with remarkable success.

For years, computer users have suspended operation and perhaps even unplugged equipment during thunderstorms [4]. Thunderstorm and lightning sensors are now made expressly for the purpose of automated thunderstorm sensing, alarm, and control. As these lightning sensors evolve and come down in price, more and more people may use them to help mitigate lightning damage and upset.

Why are engineers using these new schemes for lightning protection? Is it because protective engineering hasn't kept up with technology or because the lightning threat is worse than commonly thought?

THE LIGHTNING ENVIRONMENT

Imagine a current source that randomly connects a dozen or more times a year close to your office. When it connects, it turns on a current of around 30,000 A in a fraction of a microsecond, repeating this procedure several times in rapid succession.

In about half of these instances, hundreds and sometimes thousands of amperes continue to pass for tens to hundreds of milliseconds after the 30,000-A peak current.

Think about the metallic telephone wires, power cables, coax, plumbing, gas, and water lines that run in and out of your office. They all provide convenient conduction and induction modes that lightning currents can enter and between which voltage differences can appear.

Now imagine a nearby broadband RF noise source with a peak impulse power of 20 GW radiating into the wiring of your office, appliances, and electronic apparatus.

Just outside the walls of your building, there's the simultaneous induction of charge on exposed metal surfaces from electric-field changes of 150,000 V/m or more and EMF induction in the wiring by 300-A/m magnetic-field changes. Coincidentally, potential differences of perhaps ten to a hundred million volts are produced in the nearby soil.

Consider the various electrical and electronic subsystems in your office,

A Simple Model

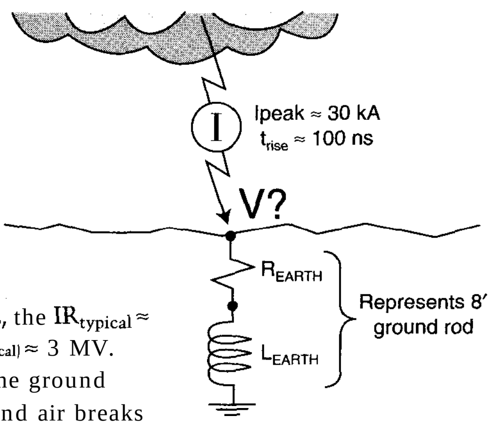

This figure offers a simple-minded electrical model treating lightning as a current source.

What is the voltage at the ground strike point at the time of the peak current?

Voltage rises are produced by IR and $L\frac{di}{dt}$ drops. If R = 10 Ω, the $IR_{typical} \approx$ 300 kV. If L = 10 μH, the $L\frac{di}{dt}_{(typical)} \approx$ 3 MV.

Since we see surface arcs at the ground strike point that are 20 m long and air breaks down at 3 MV/meter, there must be voltages at the ground strike point in excess of 60 MV to produce these long arcs. Therefore, L_{earth} may be 200 μH or so!

This simple model roughly illustrates what happens when lightning current is injected into a standard 8' ground rod. It shows that it is the surge impedance of the Earth contact-not the resistance-that matters in lightning protection.

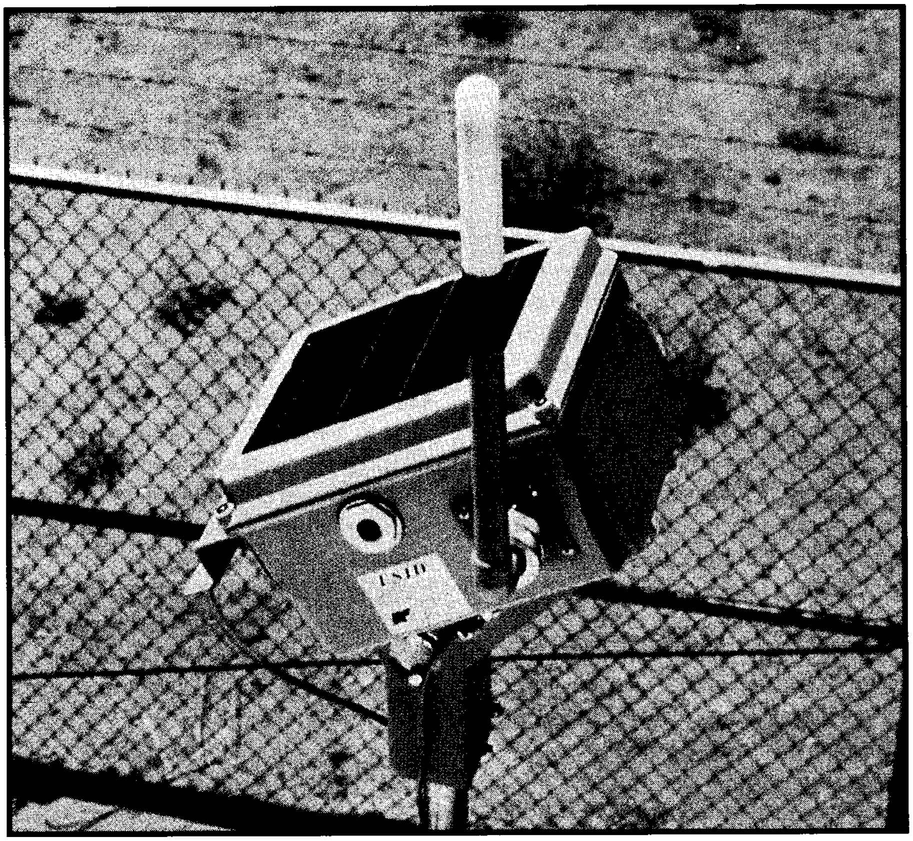

Photo 2-A *short-range, electrooptical radiation field sensor is used for preemptive lightning protection*

all the different manufacturers, and all the different contractors who install these systems. Imagine that each subsystem was designed, manufactured, and installed without understanding this threat.

Remember, it's also true that manufacturers, distributors, and installers of the electrical subsystems your office depends on actually derive significant revenue from these lightning-caused "acts of God."

Bear in mind that in central Florida, parts of Illinois, and many areas in the Southeastern United States, this threat occurs about 30 times per year per square kilometer [5].

This is the lightning environment that our electronic systems must function and survive in.

WHY THINGS FAIL

Most lightning transients appear in circuitry as common-mode voltages or currents. That is, paired or bundled conductors go collectively and momentarily to some potential quite different from normal.

For example, the line and neutral conductors of the AC power cord, normally very near ground potential, often go to the same high or low voltage relative to ground. With common insulating materials, breakdown of insulation can cause this phenomenon. It happens also when the potential of ground at some reference point is slewed impulsively from the normal value.

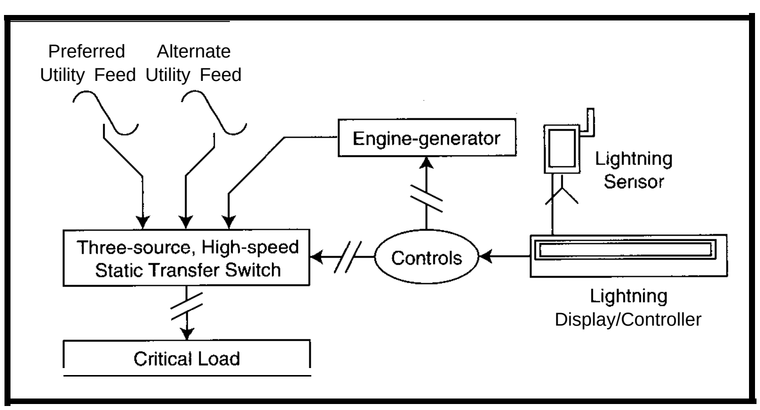

Figure I-Sarasota *Memorial Hospital uses a lightning sensor to signal the starting of the backup engine generator to avoid the complete loss of power normally associated with lightning-caused commercial power loss.*

Many lightning threats to electronic equipment involve a difference in potential between two separate inputs (such as AC power and a communications port or a parallel port and a phone line connection) that is greater than the maximum allowable range of devices and circuits in between two or more ports.

Most often, uncontrolled potential differences on the order of a few tens of volts between two separate inputs to a piece of instrumentation results in electrical overstress and failure. Good circuit designers and installers try to anticipate these threats in design and attempt to limit them.

However, most circuit designers and installers are unaware of the severity of the lightning threat and cannot control the use of the equipment. For other aspects of the lightning threat and protection problem, see the "Damage Control" sidebar.

SIMPLE SOLUTIONS

A simple solution to the lightning threat-and the seemingly intractable protection problem-is to avoid it as much as possible by anticipating thunderstorms and taking preventive action. Here's the first principle of preemptive lightning protection:

If electronic equipment is turned off, shut down, or electrically isolated from the outside world, it's much less likely to be damaged or disrupted by lightning. Complete electrical isolation of electronic equipment is as close as we can get to complete protection from lightning.

Obviously, this idea does not appeal to businesses deriving revenue from lightning-damaged equipment. Some systems cannot be shut down temporarily and so redundant power sources are used, as is common in the telecommunications industry.

In many cases, simply disconnecting from commercial power and reverting to locally derived

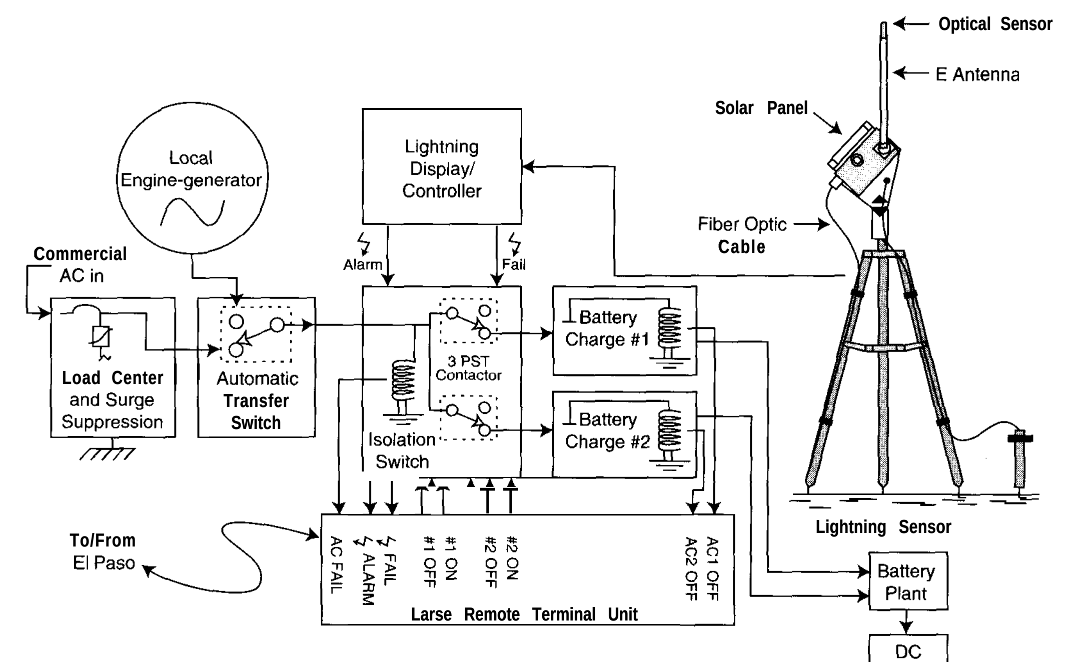

Figure 2—*El Paso Natural Gas uses a lightning sensor at remote microwave repeater sites to take the stations off of commercial power during thunderstorms.*

at a high enough rate to make for negligible leakage (between exposures) of bound charge around insulators and to the atmosphere. A state-of-the-art field mill is shown in Photo 1.

Normally, a field mill sees only the electric field (about 150 V/m) due to charged particles in the ionosphere and the random wafting of weakly electrified clouds and space charge. When a thunderstorm develops or moves overhead, the electric field at the earth increases fairly rapidly by a factor of 20 or so. This gross change is easy to detect and can be used to trigger an alarm.

power during thunderstorms removes a major offender (long commercial power lines) from the picture.

In preemptive lightning-protection schemes, your own eyes and ears can be your electronic lightning or thunderstorm sensor. Similarly, the hands that disconnect, shut down, and isolate equipment can be your own or equipment like relays, contactors, transfer switches, and the like.

Two broad classes of sensors are used for preemptive lightning protection: static and radiation field sensors, including some sensors that detect both variables. Both types of sensors can provide reliable alarms for equipment protection.

STATIC ELECTRIC FIELD SENSORS

The electric field between the base of an electrified cloud and the ground before, during, and after a lightning discharge can be monitored with a rotating-vane electrostatic fluxmeter. In the vernacular, this instrument is known as a *field mill.*

A field mill measures static and slowly changing electric fields by sensing the bound charge flowing between conducting sensor plates and ground. A grounded cover or shutter moves over the sensor plates, alternately exposing them to and shielding them from the ambient electric field.

The shutter must expose sensor plates to the atmospheric electric field

#112

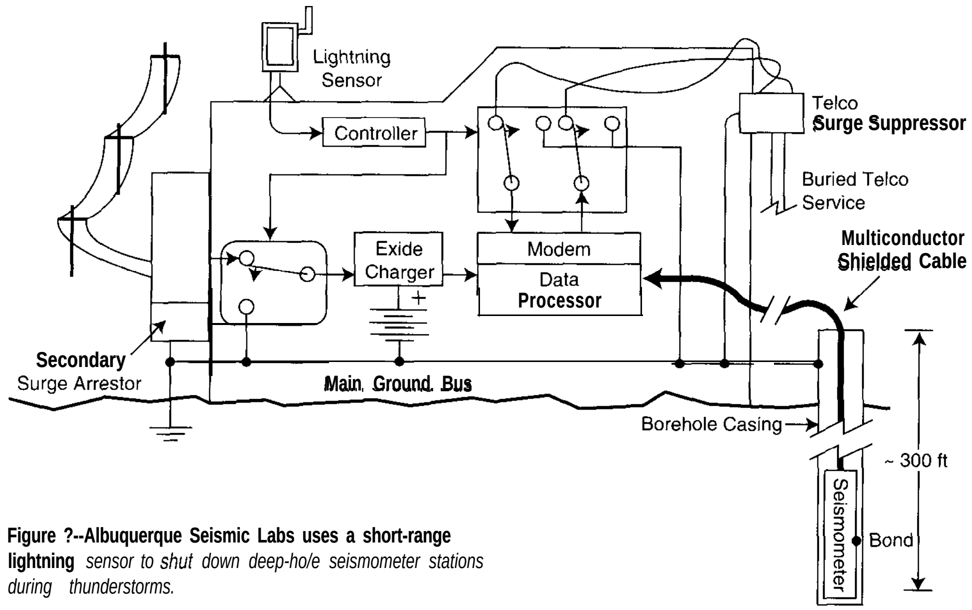

Figure ?--Albuquerque Seismic Labs uses a short-range lightning *sensor to shut down deep-ho/e seismometer stations during thunderstorms.*

Modern field mills for preemptive lightning protection include receivers and processors with relays for signaling and controlling other devices. Since commercial power often fails during a thunderstorm, any sensor used for preemptive lightning protection should operate during commercial power failures.

During a thunderstorm, the electric field is chaotic. For warning purposes, it can be ignored after an alarm or monitored for impulsive field changes to allow the counting of lightning discharges.

Typically, an alarm such as a contact closure is asserted after a thunderstorm threat is announced for some time period (approximately the lifetime of a thunderstorm). Thus, alarms don't toggle on and off many times during a stormy period. A long period of normal electric-field values signals the end of a storm.

Field mills are used in the most critical human-safety, military, and aerospace warning applications. For instance, the Kennedy Space Center has 30 field mills on the launch complex. These field mills often provide a warning before the first lightning discharge to ground.

However, field mills can be fooled. For example, a piece of Styrofoam trash under the sensor head or a dust storm can create thunderstorm-like electric fields.

RADIATION FIELD SENSORS

When a lightning discharge occurs, electromagnetic radiation across a broad spectrum is launched from the hot arc channel. Since the radiation is quite strong, almost everything acts as a lightning sensor (Allesandro Volta **stimulated** frog legs attached to a wire using radiation taken from distant lightning discharges [6]).

With proper selection of the variable(s) sensed and the bandwidth(s) employed, you can build very discriminating lightning sensors. An indiscriminant radiation field sensor usually results in false and excessive alarms and consequent failure of the protection scheme.

Any lightning sensor protecting electrical systems must be of the highest reliability. MTBFs of the sensors discussed in this article are well above 350,000 hours. Photo 2 shows an example of a modern electrooptical lightning (radiation-field) sensor designed expressly as a sensing and control element in a preemptive lightning-protection scheme.

Damage Control

To illustrate the difficult nature of the lightning-protection problem, note that lightning has to deposit 10,000 J or so into nonconducting materials of a building to do noticeable damage, but a person can be killed with an electric shock of about 100 J, and electronic equipment can be damaged with less than a joule.

So, even if lightning-protection systems are 99% effective, people in and around buildings can still be killed by residual, out-of-control effects (surface arcs, flashover, ground-step potential, vertical streamer currents, etc.).

Lightning-protection schemes must therefore be more than 99.99% effective to prevent such residuals from damaging the electrical equipment inside structures.

APPLICATION EXAMPLES

There are now hundreds of sites in the U.S. and abroad that practice some form of preemptive lightning protection. Details vary greatly from application to application, but all cases share a common sensing and control element-a short-range lightning or thunderstorm sensor.

A few notable examples of preemptive lightning protection are discussed below. They include applications in the following categories:

- hospital power plants
- microwave repeaters
- deep-hole seismometer stations
- golf-course irrigation systems
- telephone switching facilities

A preemptive lightning-protection system is in place at Sarasota Memorial Hospital, the fourth largest hospital in Florida. There, a high-speed static-transfer switch and a low-cost lightning sensor signal the starting of a 5-MW engine generator, enabling a clean transition to standby power.

Using the setup illustrated in Figure 1, the hospital avoids the 10-15-s power interruptions that otherwise occur when commercial power is lost, as a generator cranks up to speed and takes the load.

The El Paso Natural Gas Company (EPNG) operates a dense microwave-communications network for automating the natural-gas pumping and pipeline-monitoring system that covers West Texas, New Mexico, and Arizona. At some of the unmanned repeater sites, short-range lightning sensors automatically signal stations to disconnect from commercial power and revert to an on-site backup power source when thunderstorms occur within 5-10 miles. Figure 2 shows the EPNG protection scheme.

Albuquerque Seismological Laboratory designs, installs, and maintains a worldwide network of remote, unmanned, deep-hole seismometer stations. These small, complex facilities suffer severe damage when thunderstorms disrupt commercial power or when nearby lightning interacts with power, telephone, or ground systems.

Short-range lightning sensors are installed at some of the seismometer sites to sense nearby thunderstorm threats. These sensors automatically shut down and isolate the station from local telephone and electric utilities during thunderstorms, as shown in Figure 3.

Many golf courses use short-range lightning sensors to sound warning alarms for golfers and to automatically shut down computer-controlled irrigation systems. These irrigation systems are complex and expensive, microcontroller-based systems with miles of shielded and unshielded underground cable and numerous remote microcomputers.

Simply shutting down irrigation systems properly and unpowering all equipment in the field eliminates or greatly reduces lightning damage, downtime, and greens repair.

Telephone switching facilities also benefit from preemptive protection. MCI uses short-range lightning sensors at more than 30 of its large manned switch and multiswitch facilities. Facility engineers are thus alerted to thunderstorm activity around switch sites so they can make informed decisions about switching to engine generators.

ANTICIPATING THE NEXT BOLT

While we wait for improvements in lightning-protection engineering, more and more critical systems are likely to be protected with both traditional and preemptive protection schemes. On the horizon are vastly improved engineering practices. European engineers and standards groups lead the way by constructing offices with rigorous lightning-protection engineering.

The growing use of fiber-optic cable is making systems less susceptible to lightning, but watch out for that metal armor! Rocket-triggered lightning techniques test lightning-protection schemes under full-threat stresses.

Although years away from deployment, electric utilities hope to use high-speed, static (solid-state) switches in substations instead of mechanical switches. These switches will provide fast transfer of high-voltage feeders to loads so that, as lightning causes faults on one feeder, another can be brought online without interruption.

However, be wary of lightning-protection and grounding standards and codes. They are usually inadequate for solving electrical reliability problems during thunderstorms. Recall the principles you learned (or should have learned) in college that relate to high-speed circuit design, low-noise electronic design, and shielding that are based on a physical understanding of the world. The same principles can be applied to lightning problems.

And remember, if all else fails, just pull the plug. Better yet, pull the plug before all else fails.

Leon Byerley designs and builds lightning sensors and thunderstorm warning systems in his affiliation with Mission Instruments Company of Tucson, Arizona. As the proprietor of Lightning Protection Technology, he also designs lightning-protection schemes for industry and low-impedance ground systems for lightning. Leon may be reached at byerley@azstarnet.com.

REFERENCES

[1] Insurance Information Institute, Consumer News, "Summer Storms Can Be Electrifying," (press release], New York, NY, August 11, 1989.

[2] E.W. Gunther, H. Mehta, "A Survey of Distribution Power Quality-Preliminary Results," presented at IEEE Summer Power Meeting, Vancouver, BC, July 18–22, 1993.

[3] R.J. Fisher, G.H. Schnetzer, "1993 Triggered Lightning Test Program: Environments Within 20 Meters of the Lightning Channel and Small Area Temporary Protection Concepts," *Sandia Report,* SAND94-0311, UC-706, SNL, Albuquerque, NM (available from NTIS).

[4] S. Ciarcia, "Touch-Tone Interactive Monitor, A Watchdog in Every Home," *Circuit Cellar INK,* **21**, 54–64, *1991.*

[5] L.C. Byerley, K.L. Cummins, J. Tuel, D.J. Hagberg, W. Bush, "The Measurement and Use of Lightning Ground Flash Density," presented at the 1995 International Aerospace and Ground Conference on Lightning and Static Electricity, Williamsburg, VA, Sept. 26-28, 1995.

[6] B. Dibner, *Allesandro Volta and The Electric* Battery, Franklin Watts, Inc., New York, NY, 5 l-63, 1964.

SOURCES

Global Atmospherics, Inc.
Lightning Location and Protection Division
2705 E. Medina Rd., Ste. 111
Tucson, AZ 85716
(520) 741-2838

Mission Instruments Co.
5937 E. Pima St.
Tucson, AZ 85712
(520) 721-9242

I R S

404 Very Useful
405 Moderately Useful
406 Not Useful

Thick-Film Hybrid Technology

FEATURE ARTICLE

Ron Huber & John Pastre

Hybrid technology has made the shift from the simple single-layer circuit to multilayer networks of surface-mount and complex integrated circuits. Ron explains what applications are suitable for hybrid technology.

Of irst known as *hybrid circuits,* thick-film hybrid technology has made a transition from its beginnings as a simple single-conductor-layer circuit interconnecting discrete semiconductor-chip capacitors and printed planar resistors. It is now capable of multilayer networks interconnecting a wide range of surface-mount devices as well as complex integrated circuits (see Photo 1).

As a mature packaging technology, it continues to evolve, using modern technology to keep up with the ever-expanding needs of electronics and digital industry. It has expanded to meet the needs of many new electronics and digital applications, especially applications requiring reliable miniaturized networks.

Thick-film hybrid technology is certainly not a technology for all applications But, after this overview, you should have a better appreciation for its niche.

We'll review what a thick-film hybrid circuit is and then look at what applications it suits and how it compares to other packaging approaches.

HYBRID EVOLUTION

The basic thick-film process is an outgrowth of silkscreen technology, which is many centuries old. The fabrication of electrical circuitry using screen-printing techniques gained impetus in the 1960s when Dupont developed highly reliable compositions for screenable resistors.

This development enabled low-cost fabrication of passive circuit elements as part of the same planar structure as the conductors. With a snowballing effect, this capability stimulated the development of improved screenable conductor materials based on systems using noble metals such as gold, platinum, and palladium.

The availability of these improved conductors and wide range of resistor pastes marked the true beginning of the thick-film hybrid circuit. Much of the drive to develop these hybrids was for size reduction-in many cases, the hybrid was only 25% of the area of the discrete approach.

The circuit involves a network of screened conductive paths interconnecting screened resistors or capacitors and semiconductor devices on a rigid ceramic substrate. Usually a dielectric material is screened for conductor crossovers or a multilayer interconnects. Figure 1 illustrates the composition of a hybrid circuit.

THICK-FILM HYBRID PROCESS

The thick-film hybrid fabrication process is strictly an additive one and involves no photo etching. The process consists of screen printing, air drying, and firing layers of conductors, dielectrics, and resistors sequentially on a ceramic-based substrate.

The substrate material most frequently used is alumina (Al_2O_3), which is 92-99% pure depending on the surface smoothness required. Ceramic substrates may be purchased to finished-circuit size or by laser-scored multiple substrates that are snapped apart after batch processing is done.

Artwork is generated for the conductor, via, dielectric, and resistor layers using CAD tools with built-in design rules to implement layouts for most applications. The artwork is transferred to a photosensitive emulsion on fine-mesh wire screens to form a negative pattern that allows precise deposition of viscous conductor, dielectric, and resistor pastes.

SCREEN PRINT AND FIRING

Hybrid circuitry is fabricated by screening and firing each layer sequentially through its respective patterned mesh screens. In multilayer structures, vias are typically 8–1 5-mil openings in the dielectric layers that are filled with conductor material in the subsequent conductor printing layer.

If resistors are needed, they're usually printed as the last step to maintain resistive control and stability. A 10% tolerance may be achieved with a resistor well into the megohm range.

Laser trimming to less than I % is readily achieved. Recent developments for printed resistors mean that they can be buried within the interconnect structure and still maintain good characteristics.

The various thick-film paste materials used in the printing process have active elements held together in a viscous form by volatile binders that are later burned away during the high-temperature furnace-firing process.

The most widely used active elements in thick-film cermet conductors are conductive particles of either silver, palladium, gold, or platinum, depending on the resistivity and processing characteristics required. Functional elements in dielectrics are glass, ceramic, or a combination of both. Resistor pastes contain ruthenium oxides or other ruthenates formulated to achieve a range of resistivity.

The thick-film patterns are fired after each printed level, usually in a temperature-controlled, multiphase belt-fed furnace. A temperature profile allows a buildup to about 850°C. This buildup causes centering or welding of the discrete particles of the functional fillers in the conductors, dielectrics, and resistors.

ACTIVE DEVICE Al-t-ACHMENT

The thick-film interconnect structure described earlier provides an excellent medium for mounting a wide variety of surface-mount devices. Such devices include miniaturized surface-mount packages (QSOP, VSOP), leadless chip carriers (LCC), flip chip (FCP), TAB, chip and wire, or any surface-mount device (SMD) down to IO-mil pitch mounting.

Platinum-gold surface pads and screenable solder pastes enhance the solderability of surface-mount packages, while gold pads offer reliable automatic wire bonding of unencapsulated chips. Differences in chip height are accommodated by high-speed autofocus, automatic wire bonders.

The ability to mix the two attachment processes on the same hybrid surface yields additional versatility. Uncased chips and wire bonding permit networks of the highest density by eliminating the package overhead area.

As an example, a 28-lead DIP is 1.4" x 0.6" and requires a bonding area of 0.85 in2 If the same chip in the package is chip-and-wire surface mounted, the area required is only 0.2" x 0.2" or 0.04 in^2 The DIP requires over 20 times the area of the uncased die.

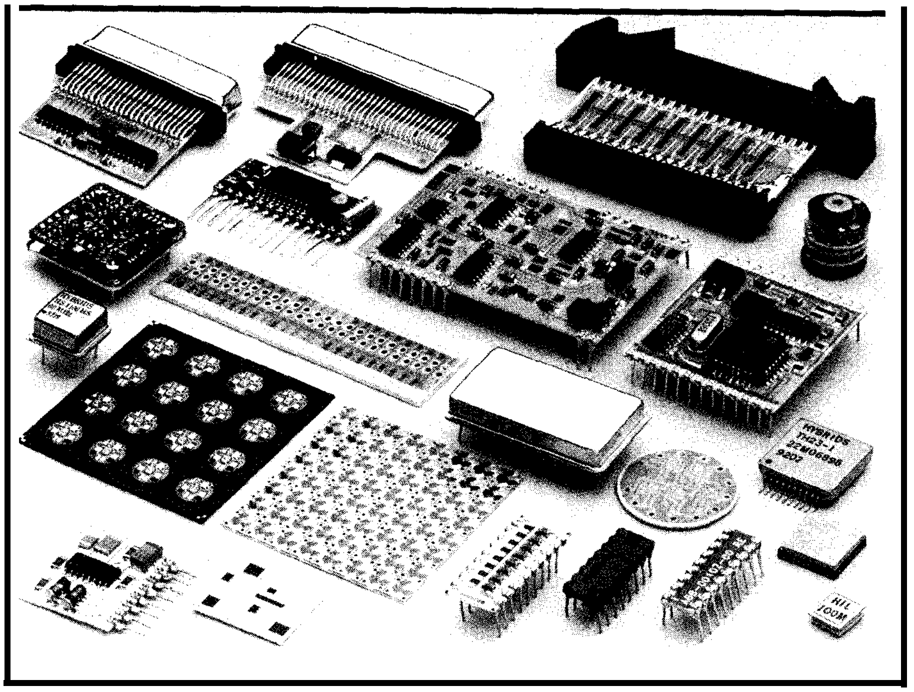

Photo *I-Typical hybrids use various substrates and packaging. Custom circuits have been designed for custom applications.*

MULTILAYER TECHNOLOGY

New-generation electronics and digital applications demand more functionality, higher operating speeds, greater power densities, improved reliability, and of course, lower cost. That's a tough set of requirements for any packaging approach to satisfy, and at this time, perhaps none can meet them all. However, thick-film technology has, to some degree, advanced to address these unique requirements.

Increased functionality requires a greater number of line nets to interconnect more complex semiconductor devices. A typical digital board or module application today may require 1000 interconnect nets as opposed to 500 in the past. When dealing with 300-pin devices in the future, 4000-line nets may be required.

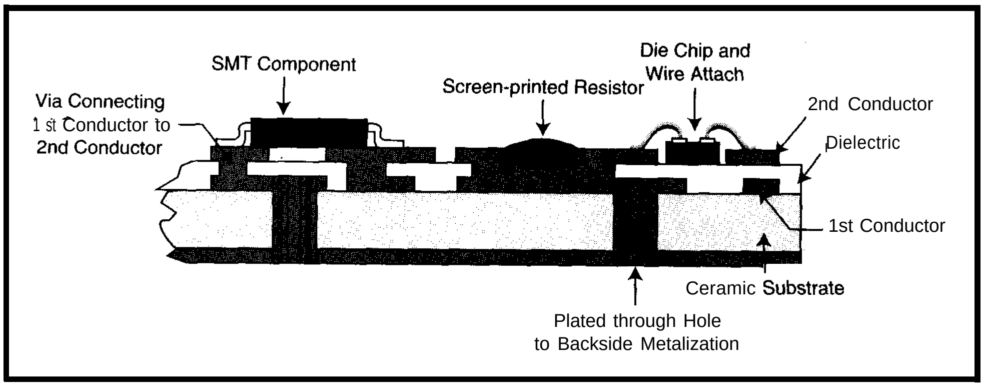

Figure 1—*A thick-film hybrid with printed resistors uses both surface-mount and chip and wire technology.*

All this adds up to a requirement for higher-density interconnects and more signal, ground, and voltage layers. Thick-film technology has met this need by developing higher-density interconnects down to 4-mil lines on 10-mil pitch and 10-mil vias.

The capability for buried interlayer vias enhances the interconnect and device-attachment density by not requiring drilled plated-through holes. Also, reliable vias established through the base ceramic substrate and around the edges enable multilayer interconnects on both sides of the substrate. These developments mean devices can be mounted on both sides of the substrate, increasing functionality by up to a factor of four.

We know of one successful Multi-Chip Module (MCM) application of thick-film multilayer technology where 12 complex digital and memory chips were attached and interconnected on a 2" x 2" thick-film substrate and further mounted in a 240-lead low-profile flat pack. This package not only saved a great deal of surface area, it also potentially improved reliability by reducing I/O counts from over 3000 to 240 by reducing PC board connector interfaces.

APPLICATION DECISIONS

Where does the thick-film hybrid network best fit? Where should it be used? Obviously, it's not meant for all applications, but where it's best?

In reality, hybrids are in a broad range of applications in the consumer, military, and industrial sectors. At Hybrids International, we've found a demand for hybrids of various complexities where miniaturization and ruggedness are needed.

They're useful in:

- communications
- defense
- avionics
- automotive
- instrumentation
- computers

Within these application areas, these projects use hybrids as a key

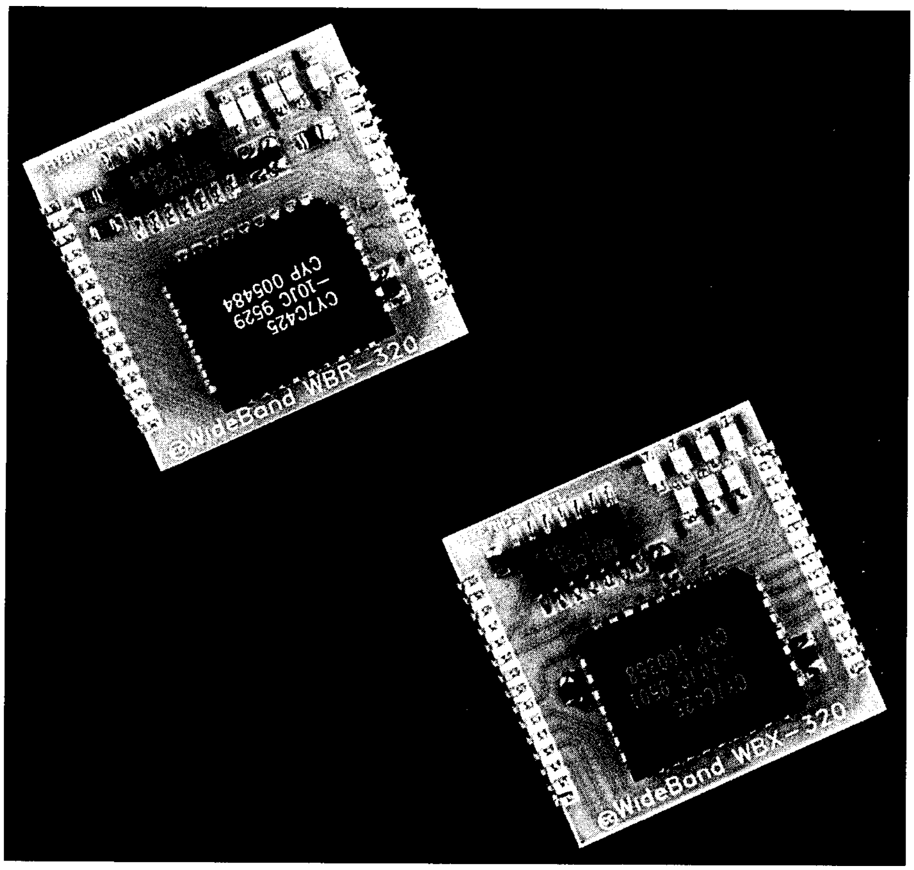

Photo *2-Custom* **transmitter and receiver circuits using thick-film** *technology have surface-mount devices on both sides.*

packaging approach or to support other circuitry:

- mobile phones and pagers
- Global Positioning System (GPS)
- engineering workstations
- portable instruments
- portable/laptop computers
- modems, fax machines, copiers
- fiber-optic communications circuitry

Obviously, the need for size reduction is key. As noted, with the use of miniature surface-mount components and chip-and-wire devices, a 2–4-times reduction in size may be realized over widely used packaging such as epoxy boards with DIPs, PGAs, and discretes. Photo 2 highlights the degree of miniaturization that can be achieved.

The use of surface-mount devices on epoxy boards has increased their potential density. However, thick-film ceramic-based technology still provides the greatest density and the best cost-for-density factor.

Military applications provided much of the drive in developing thick-

Comparison Parameter	Conv. EGL PC Boards	Thick-Film Hybrids	Thin-Film Hybrids	Cofired Ceramics
Conductors				
line width	8-10 mil	4-8 mil	0.5-l mil	8-l 0 mil
material	cu	Au,Cu,Pt-Au	Cu,Al	W,Mo
Dielectrics				
material	epoxy-glass	ceramic/glass	polyimide	ceramic
const. K	4.8 (med)	8.5 (med-hi)	3.5 (low)	9.5 (hi)
Resistors	discrete	printed	deposit	discrete
Component attachment	DIP,PGA, leaded SMD, no chips	SMD,BGA,FCP LCC,TAB, chips	SMD,BGA,FCP LCC,TAB, chips	SMD,BGA,FCP LCC,TAB, chips
Pkg. density	med	med-hi	hi	med-hi
Power density	fair	good	good	very good
Electrical performance	good med K long lines	good med-hi K short lines	very good low K short lines	fair hi K short lines
App. flexibility	good	good	fair	fair
Relative cost	low	med	hi	med-hi
Capital equipment cost	low	med	hi	hi
Custom design vendors	many	many	limited	limited

Table l-Various *substrates and technologies are used in* hybrid-circuit manufacturing, *and depend on the application requirements.*

film hybrid technology, stimulating hybrid-circuit growth in the 1970s and 1980s. Military equipment required miniaturized, rugged, and reliable electronic networks.

Hybrids provided reliability enhancements through a rugged thick-film interconnect structure composed of fairly inert materials such as ceramic, glass, and gold. These materials resist degradation and failures in hostile military environments.

The reduction in interconnection interfaces, when using uncased devices, and the many gold-to-gold connections have demonstrated reliability. The use of hybrids in military applications increases the ruggedness and reliability. These factors have led to hybrids being considered for applications like mobile phones, pagers, portable equipment, and smart sensors.

Smart sensors in particular may benefit because of the need to locate the data-acquisition, reactive-control, and first-level-communications hardware at the point of contact.

A high-density and reliable electronics-packaging approach is required in many of the severe sensing-point environments. Automobiles, for instance, provide a likely environment.

POWER APPLICATIONS

Applications generating high power benefit from the thermal-handling capabilities of ceramic-based thick-film networks. Such applications might include voltage regulators, voltage converters, switching power supplies, and other high-power circuits.

High-power devices such as power diodes, transistors, and SCRs mount directly to the ceramic-based interconnect structure and are automatically heatsinked to a high thermal conductive structure. As well, they serve as the interconnect media.

The alumina ceramic substrate has a thermal conductance (K) of approximately 0.8 W/°C-in. This K is about 10% of the conductance of pure copper but about 10 times better than most epoxy-glass PC materials.

A beryllia substrate can be used in designs where extremely high power handling is required. Beryllia has a K constant of 5.5 W/°C-in., approximately 55% that of copper. A 200 x 200-mil power chip mounted on beryllia has a thermal resistance (R) of approximately 0.4°C/W. Therefore, with $P = T/R$, at a 20°C allowable temperature rise, this chip could dissipate 50 W continuous. In a normal mounting configuration, the result is a power density of approximately 300 W/in2

In some MCM applications, chip densities are high and power densities may rise to 2530 W/in2 For these applications, the thermal capabilities of a ceramic-based interconnect are needed to manage the thermal issues of temperature rise and stability.

PACKAGING TECHNOLOGY

We're sure it won't be disputed that the king of electronics packaging is still the epoxy-glass laminate-based technology (EGL), and for good reason.

It's versatile and can accommodate through-hole devices such as DIPs and PGAs as well as newer surface-mount elements down to 20-mil lead pitch. Cost is moderate, and there are many PC board vendors who supply the interconnects for a design that can be assembled and tested in-house.

EGL technology does, however, start to run aground when applications require miniaturization or a more rugged environment. Situations may arise when an IC isn't available to perform a required function and must be designed in a discrete configuration. Designing with conventional discretes is inefficient in that it takes up expensive board space.

Hybrids are a good solution here. A hybrid can be designed, packaged, and mounted on the board as if it were the IC, usually with a great savings of board space. At the other end of the scale, where the entire design function must be miniaturized to meet a form fit or a performance window, hybrid technologies start to shine again.

The most widely used hybrid technologies that compete for this miniaturization industry are thick film, thin film, and what is known as *cofired ceramic.* All three have been used extensively and have various advantages depending on the application.

Table 1 provides an overview of the four packaging technologies. There are a great many more parameters to compare. However, those listed in Table 1 are considered key and raise the greatest concerns.

The comparisons in Table 1 lead to these simplified conclusions:

- EGL technology has the greatest flexibility when using present widely used device packages. It has good performance and cost factor with many readily available sources. It appears to be more limited at high power densities and certainly its packaging density is limited by the inability to use uncased chips or SMD at lead pitches below 10 mils.
- thick-film hybrid technology has excellent flexibility when using surface-mount and uncased devices. It appears to have the best cost-for-density factor and has many vendor sources for custom developments.
- thin-film networks have good flexibility when using surface-mount devices and chips. This technology yields the best density and performance for MCM, but at higher cost. The number of custom sources is more limited.
- cofired ceramics have similar attributes to thick film. However, performance is slightly degraded, cost is higher, and custom vendor sources are much more limited.

CONCLUSIONS

In this overview of thick-film hybrid technology, we've established a baseline of its characteristics and capabilities and shown how it fits into applications from a packaging standpoint. Thick-film hybrids have found a niche in the miniaturization and ruggedness sector over the years.

With a newer generation of applications driving the need for higher performance and functionality per unit of space, it appears that thick-film networks will just keep fitting better.

As we see it, hybrid technology will be able to satisfy the packaging requirements of a wide range of the newer electronics and digital applications. ❑

John Pastre is currently a consulting engineer in electronics, components, and packaging. He spent a number of years managing the hybrid design group at Unisys and as a technical specialist addressing MCM developments at MCC in Austin, TX. He has authored six articles on hybrid circuits and MCM development.

Ron Huber is currently the sales manager at Hybrids International. He has spent the last 25 years in the thick-film hybrids industry at companies including CTS, Marconi, and Sperry Corp. (Unisys) in both sales and engineering positions. Both authors may be reached at hybrids@idir. net.

SOURCE

Hybrids International, Ltd.
707 N. Lindenwood Dr.
Olathe, KS 66062
(9 13) 764-6400
Fax: (913) 764-6409

I R S

407 Very Useful
408 Moderately Useful
409 Not Useful

Microvolt-Level Data Acquisition

FEATURE ARTICLE

Bob Perrin

Microvolt design is no easy task, especially for the uninitiated. Here, Bob exposes the problems that come with microvolt design and explores specific solutions to those problems.

Design of microvolt-resolution data-acquisition systems seems a daunting task for the uninitiated. But, take heart. Microvolt resolution is obtainable by all who truly fear their boss.

The first thing to remember is that all traditional ideas of resistance, reactance, gain, and so on are still valid at microvolt levels. Even though microvolts are somewhat unpredictable, triboelectric, electrochemical, and dielectric effects are still manageable.

Submicrovolt design is where the real challenge lies. Inductive and capacitive coupling are serious problems. Small leakage and bias currents can produce significant voltage errors. Op-amp offset voltages are huge compared to the signal levels. Resistive loading can devastate sensor outputs.

In this article, I focus on practical problems and specific solutions. Check the references for in-depth discussions of theory and mathematics.

SYSTEM OVERVIEW

The concepts outlined in this article are tailored to low-frequency, high-precision data acquisition. Figure **1** shows the circuit topology referenced in this article.

The I/O block constitutes the instrument's interface to the human world and includes relay and motor drivers. Generally, this block has low activity, but high current requirements.

The microprocessor block contains all the elements typical to an embedded controller: memory, glue logic, and the processor. This subsystem is characterized by moderate current requirements and large digital-switching noise.

However, the low-noise amplifier, A/D converter, and precision reference blocks are where the exciting action takes place. The sensor block contains the necessary hardware to convert the sensor's electrical output to a voltage.

Common conversion schemes are outlined in Figure 2. Resistive sensors have an output resistance that varies proportionally to an input phenomenon (e.g., temperature). A sensor's resistive output can be converted to voltage by an R-to-V converter. Although self-heating can be a problem with these sensors, a short low-current pulse, applied only for the duration of the measurement, reduces the self-heating effect.

A current-mode sensor generates a current proportional to an input phenomenon (e.g., the photon flux incident on a photodiode's die as pictured in Figure 2b). A transimpedance amplifier is a commonly used tool for current-to-voltage conversion.

LAB TECHNIQUE

Careful lab technique is required with microvolt-level signals. The best technique involves an elegant mixture of common sense and a solid understanding of physics.

To start with, you need to bear in mind the difference between precision and accuracy. Don't be fooled by the seven-and-a-half digits on a $2200 DMM. Although it may be precise, it may not be accurate. The sidebar "Important Distinctions" offers an extended discussion and example of how these terms differ.

All test probe connections form a thermocouple. Keep the test probes short, shielded, and away from EM1 sources. The test instrument ground must be connected to the point on the PCB where the measured signal is referenced.

There are special tips for most scope probes that let the user connect to ground near the probe tip and eliminate the ground pigtail. This feature requires PCBs designed with special test points. While desirable, the use of special tips is not usually required for rough tracking of noise sources.

The scope is grounded through the AC wall outlet and can cause troublesome ground currents. Keep all test equipment on the same circuit, preferably on the same power strip. And, since the human body acts as an antenna, keep your hands away from signals being measured.

Unused electronic equipment should be kept away from the work area to prevent EMI. Although computers are obvious offenders, desktop telephones, pocket cell phones, radios, keyboards, switching power supplies, soldering irons, and electric fans can also cause problems.

Thermal gradients create thermal EMFs. Don't attempt precision measurements on circuits that have just been soldered, washed, turned on, or handled. Let the circuit reach thermal equilibrium before attempting measurements. Place a shoe box over the PCB under examination to reduce troublesome air currents.

Many high-precision test instruments require warm-up time. Observe manufactures' recommendations.

All voltage measurements load the circuit being measured. Don't try to measure the voltage on a S&H (sample and hold) capacitor and expect the S&H circuit to continue to function.

A solid battery of tools is indispensable for microvolt-level work. The sidebar "Tools and Toys-A Christmas Wish List" presents several useful tools.

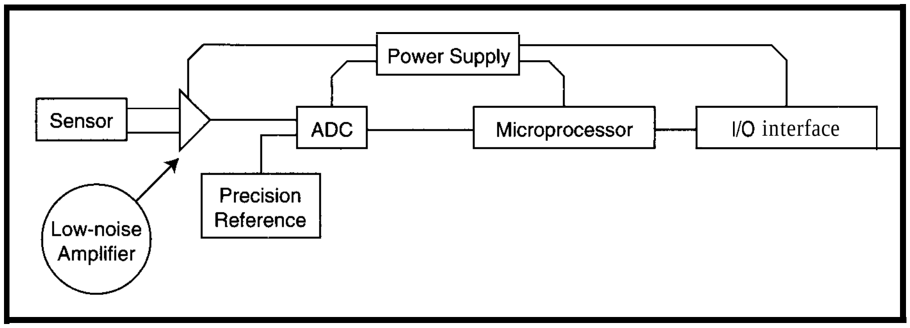

Figure 1-A *typical* **data-acquisition system** *can be broken up info several functional blocks. Each block can then b developed* **independently.**

NOISE

Electronic noise is terribly bothersome. Tiny noise sources, routinely ignored in macrovolt circuit design, tend to obliterate microvolt information signals.

Noise sources are divided into two categories. Noise generated within the circuit components is intrinsic noise while noise coupled from outside the circuit is considered external. External noise includes any noise not associated with random movement of charge carriers inside circuit components.

Often external noise is quite sizable compared to the information signal. External noise is generally not random. Its frequencies are identifiable and treatable. The sidebar "A Maze of EMI" presents a number of real-life examples.

For external noise to show up in the low-noise analog, two elements are required: a noise source and a coupling mechanism. To eliminate unwanted noise, the designer must eliminate one of these elements.

In a lab setting, it is often possible to physically remove the offending noise source, such as a nearby telephone. However, once the low-noise instrument is in the customer's hands, the engineer has no control of the external environment. The prudent engineer works therefore to eliminate available coupling mechanisms.

External noise can often be reduced to tolerable levels through classic techniques. A brief discussion of several common tools for noise reduction follows.

SHIELDING

Cabling is very susceptible to induced noise. Of the two types of cable shielding-braid and foil-a nice tightly woven braid offers a superior shield.

For microvolt signals, both a foil and a braid shield are recommended. Many cable manufacturers offer cables with both foil and braid shields.

Triax is a good cable for preserving microvolt signals. Triax consists of an inner solid core, an intermediate braid conductor, and an outer braid shield. Triax is much like shielded coax.

Triax is fairly expensive as are triax connectors. Often a good foil-plus-braid shielded conventional cable suffices. A conventional cable offers the ability to maintain twisted pairs internally, thus maximizing the effectiveness of the common-mode-rejection (CMR) characteristics of the low-noise differential amplifier.

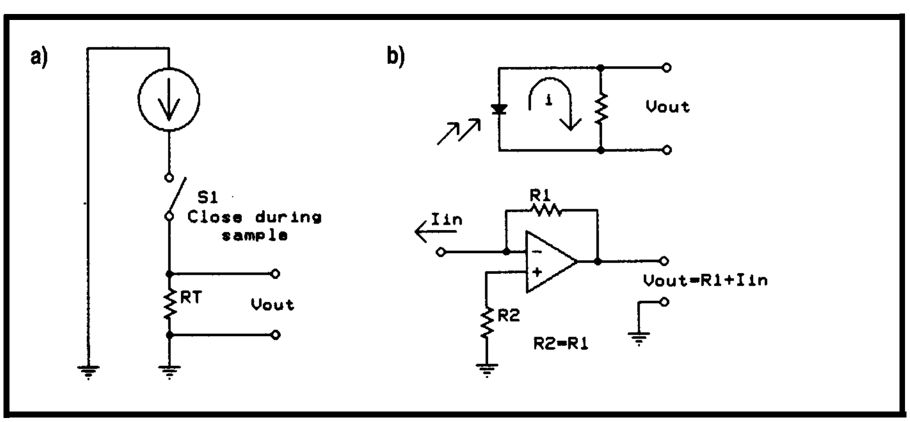

Figure 2—*Often* **an** *engineer must convert a* **resistance** *or* **current** *to a voltage.* **Three** *simple converters* **are shown.**

POWER-SUPPLY DECOUPLING

Place decoupling caps on all ICs. A 10-µF and a 0.1-µF cap are commonly recommended by analog IC manufactures.

One advantageous product is the Q-2000 line of caps from Circuit Components. These flat, skinny capacitors are designed to be installed under digital DIP ICs. Through geometry, the lead inductance is'minimized. Manufacturers gladly supply specific data.

Decoupling the digital ICs with a Q-2000 cap and conventional capacitor squeezes out as much digital-switching noise as possible. A Q-2000 cap, a 0.1-µF 1206 SMT cap, and a 10-µF radial tantalum are recommended for each digital IC.

POWERING SUBSYSTEMS

One technique that reduces noise originating from the microprocessor section is to simply shut down the processor during an A/D conversion. However, three serious pitfalls exist:

- unpowered circuits may unpredictably load powered circuits, causing sporadic erroneous signals
- powering up and effecting a reset may be very slow (i.e., hundreds of milliseconds)
- the microprocessor loses its brain and on reboot must be reminded who and where it is.

All of these problems are solvable.

Many processors have a sleep or idle mode. Often the master clock is taken down when the processor enters one of these low-power modes, which is just as good as shutting off power. Frequently, there are inherent facilities for maintaining critical stack and register information when employing sleep or idle modes (i.e., the processor does not lose its brain).

POWER

Using multiple regulators is a quick way to clean up noise. Provide the digital power with a separate regulator than the one supplying analog power. The MAX667 5-V, adjustable, low-dropout voltage regulator is an 8-pin bug and has many possibilities.

At the minimum, use two separate grounds: a noisy digital ground and a quiet analog ground. The analog ground system should be implemented with a dedicated ground plane.

These two ground systems must eventually be referenced. Intuition suggests connecting the independent ground systems at the power supply.

However, experimentation often yields a better, sometimes counter-intuitive, location. A Crystal Semiconductor applications engineer recommends connecting the grounds at the A/D converter. The Crystal Semiconductor databook has several terrific application notes [2].

If high-current drivers exist in the I/O interface, a third ground should be employed.

COMMON MODE REJECTION (CMR)

CMR is an analog engineer's best friend. If the leads from the sensor can be arranged so that induced noise is common to both leads, then a low-noise differential amplifier can attenuate many decibels of unwanted external noise.

Twisted pairs are best for keeping induced noise common mode. Shielded cables with internal twisted pairs are available. This type of cable is recommended for low-noise transducer signals.

Carefully examine graphs (not marketing hype figures) of the CMR characteristics of amplifiers under consideration. CMR generally rolls off with increasing frequency. A CMR of 120 dB isn't useful if it rolls off to 10 dB at 60 Hz.

Important Distinctions

Precision, accuracy, and resolution are separate terms with specific definitions. Precision refers to the repeatability of a measurement, accuracy to the closeness of the measurement to the true physical value, and resolution to the smallest detectable change in a measured quantity.

A homebrew DMM illustrates my point. I measured a 1-V ideal source five times using a 16-bit ADC referenced to a 10-V ideal reference. If you look at Table i, you'll see the ADC codes in column 2 and the instrument's display in column 3. The display precision is tabulated with a 0.1 -µV precision.

The measurement precision is:

$$\begin{aligned}\text{precision} &= \text{maximum} - \text{minimum}\\ &= 1.1062622 - 1.0910034\ \text{V}\\ &= 0.0152588\ \text{V}\end{aligned}$$

The converter resolution is:

$$\text{resolution} = \frac{\text{reference voltage}}{\text{A/D converter bits}} = \frac{10\ \text{V}}{2^{16}} = 152.587\ \mu\text{V}$$

The measurement accuracy is:

$$\begin{aligned}\text{accuracy} &= \text{average display V} \times \text{resolution} - \text{true value}\\ &= 1.0988172 - 1.0\ \text{V}\\ &= 0.1\ \text{v}\end{aligned}$$

This is a 10% error from the true value.

Resolution is determined by the choice of converter and sampling techniques. It can be increased by oversampling and applying DSP techniques.

Display precision is determined by the number of digits carried mathematically in software. Measurement precision is determined by system noise.

Accuracy is affected by noise, offset errors, reference errors, thermal EMFs, sensor errors, and so on.

In this example, we have a system with excellent resolution, mediocre precision, and terrible accuracy. This system probably suffers from moderate noise levels (degrading precision) and a big amplifier offset or thermal EMF (degrading accuracy).

Keithley Instruments' ***Low-Level Measurements Handbook*** offers more detailed discussion [1].

Trial	ADC Codes	Displayed Volts ADC x 1 0/2^{16}
1	7209	1.1000061
2	7150	1.0910034
3	7250	1.1062622
4	7220	1.1016845
5	7180	1.0951298

Table i-A home-brew *DMM* was used *to measure a 1-V source.* Five trials *were run.*

A Maze of EMI

Some true-life EM1 encounters are recounted here:

- 60-Hz line noise on sensor wires is EM1 inductively coupled from power cords or building electrical circuits.
- 60-Hz noise (and/or harmonics) from the rectifiers in the DC power supply can be conducted on the amplifier power rails. These annoying harmonics can also be coupled into low-noise sections.
- 120-Hz noise can be optically coupled into exposed silicon. Photodetectors, optoisolators, and unshielded thermopiles are susceptible. The intensity of artificial lighting varies with power, not voltage frequency.
- A charge-pump voltage converter does not work with microvolt circuits. The gulps of current swallowed by the charge pump create an impressive amount of EMI. This noise is both easily radiated and conducted.
- Weird but identifiable megahertz frequencies are generally traceable to a microprocessor bus. The offending microprocessor may or may not reside in the circuit under development. Nearby computers are frequently the guilty perpetrators.
- The microprocessor crystal frequency can be broadcast or conducted readily. One common solution to this problem is to sleep or idle the processor while the A/D converter is sampling. The A/D converter can wake the processor after the conversion.
- Some temperature-controlled soldering stations produce monstrous EM1 pulses. To regulate tip temperature, soldering irons switch large currents to a heating element. This switching broadcasts noise pulses. Low-frequency (seemingly intermittent) noise spikes can sometimes be traced to a nearby soldering station.
- Mechanical vibration of sensor wires produce difficult-to-trace noise. This vibration can be a friction-induced voltage caused by the triboelectric effect. More commonly, the noise is due to inductive coupling from a magnetic field. Secure your cables well!
- Air currents creating thermal gradients in the low-noise circuit are difficult. Fans should be avoided in enclosures with low-noise analog. A thermal mass (0.125" copper plate] can be placed over the low-noise components to reduce thermal gradients.
- Some pocket cell phones cause low-noise circuits to go bonkers. Just remove the batteries from your trusty pocket cell phone.

Instrumentation amplifiers (IA) provide maximum CMR and are easily understood. The classic three-op-amp topology is presented in every modern text on amplifier design, though several difficulties haunt this configuration.

Exact resistor matching (0.1% are indicated) is required to obtain maximum CMR. Trimming is required for individual op-amp offsets. The more components, the more susceptible to temperature gradients the physical circuit becomes. These difficulties are manageable with careful circuit design, component selection, and PCB layout.

The cheater's way around the idiosyncrasies of the three-op-amp topology is the single-chip IA. Several manufactures offer single chip IAs—most notable of these is Burr-Brown.

The Burr-Brown INA103 has some of the lowest noise specs available. However, the part dissipates enough heat that small air currents cause serious offset voltage fluctuations. If this device is placed in a sealed enclosure and allowed to reach thermal equilib-

Tools & Toys-A Christmas Wish List

Proper lab technique and a good array of tools are two elements facilitating a short design cycle. A few basic instruments are extremely helpful:

- a DSO with a good FFT package-the FFT package lets the scope be used as a (limited) spectrum analyzer. This instrument is useful in identifying and isolating noise sources.

  Many people have experienced problems with the FFT packages provided with lower-priced DSOs. Don't even bother with the cheap combination scopes. Le-Croy scopes are the best DSOs made. The LeCroy 9300 series scopes are excellent values.
- an analog scope-l mV per division is a preferable low end for the lx probe. Tektronix, HP, Philips, Kenwood, B&K all make adequate low-end analog scopes.

  Cheap analog scopes don't have repeatable or stable trigger circuits and may have jitter in the input pass-band. These features are tolerable, however, since the analog scope is used to look at the magnitude of fuzzy lines. It gives quick estimates of noise levels.
- an instrumentation-grade preamp-This device can be placed on the front end of a scope or meter to increase sensitivity. A variable gain from 1 to 1000 is handy.

  The Stanford Research SR560 provides best performance for the cost. It has configurable high- and low-pass filters, which help evaluate the effects of band-limiting noise.
- a precision voltmeter-Submicrovolt testing may demand a nanovoltmeter.

  The Keithley 2001 and 2002 are good DMMs. Both have option slots that permit a nanovolt preamp to be installed. Input impedance, source impedance, resolution, and sampling rate are all relevant issues when selecting a meter. Keithley Instruments' *Low-Level Measurements* [1] has an excellent discussion on selecting a meter.
- a sealed and grounded 2.5' x 2' x 4' metal box-This miniscreen room can shield the circuit under test, test instrumentation, and probes from EMI. Inside the enclosure, air currents are minimized, which also minimizes thermal transients. Thermal conditions affect offset voltages and thermal EMF generation.
- a thermocouple is an excellent low-impedance (10 Ω) voltage source. A large thermal mass dampens the effects of air currents. An aluminum block 2" cubed, with a 0.060" hole drilled to the center of mass, is an excellent (and inexpensive) thermal mass. The thermocouple is inserted into the center of the cube, and secured with five-minute epoxy. The aluminum block is then covered (all 6 sides) with 1" packing foam. Keep thermocouple wires as short as possible and tightly twisted.

  Minimize air currents at the cold junction to further enhance stability.

rium, it might work nicely. It holds promise for low-noise designs.

The Burr-Brown INA118 is another IA and is also a so-called micropower device. The noise specs are significantly higher than the INA103, but the power dissipated is only a fraction of that dissipated by the INA103.

The INA118 is not susceptible to small air currents, and as an added bonus, is packaged in an 8-pin bug. It has been used successfully in submicrovolt designs.

POWER-SUPPLY REJECTION (PSR)

PSR is the ability of an amplifier to decouple noise on the supply rails from the amplifier's output, which is extremely useful. Supply noise is difficult to eliminate entirely.

Generally, PSR rolls off with increasing frequency. Some devices have notches in their PSR response that occur at most unfortunate frequencies. For example, the MAX406 has a dip in PSR at 4 kHz.

When selecting amplifiers (and references), look carefully at PSR graphs. Again, examine graphs, not marketing hype specs.

SELECTIVE FILTERING

Filtering noise in a system is generally more of a pain than preventing noise from being coupled into the system. Shielding, CMR, intelligent deployment of power supplies, and a solid grounding scheme should be employed first. Analog filtering should be done as a last resort.

INTRINSIC NOISE

Intrinsic noise is generated internally to the circuit components and can be caused by many things. The sidebar "Intrinsic Noise Sources" expands on some of the causes and characteristics of intrinsic noise.

Many factors affect the motion of charge carriers. Thermal excitation, probability of a carrier hopping a junction, recombination times, and die contamination all affect carrier motion.

Since intrinsic noise is by nature random and unavoidable, certain design techniques minimize intrinsic noise.

Many texts place high value on maximizing SNR initially by using the Maximum Power Transfer Theorem (MPTT). Matching source impedance to input impedance ensures maximum signal power is transferred.

However, using MPTT often introduces loading errors. If the sensor impedance is known exactly, the loading error can be computed and corrected. Most often the source impedance is not known exactly. MPTT seems to be most useful between amplifier stages and should be used whenever possible.

To preserve SNR, place maximum gain in the first amplifier stage. Lift the itty-bitty sensor signals out of the sea of intrinsic system noise.

Once the small signals are up to 100-µV levels, the intrinsic noise asso-

Intrinsic Noise Sources

Intrinsic noise has many causes and is not fully understood. Generally, noise is described by its power spectrum.

White and pink noise refer to specific power spectrums. White noise has power that is the same in any equal-sized frequency range regardless of where that range is. For example, white noise has the same power in the 1–2-Hz band as the 100–101-Hz band.

The power spectrum of pink noise varies $1/f^2$. Just as very little power exists at high frequencies, very large power exists at low frequencies.

Different mechanisms generate noise in different materials. Several types of noise sources are commonly encountered:

- Avalanche noise—PN junctions in reverse bias exhibit avalanche noise, which is white noise. Avalanche noise is of greater power than shot noise, but caused by similar mechanisms.
- Burst noise (or popcorn noise)-This noise is generated only in semiconductors and appears to be caused by metallic die contamination. The power varies as $1/f^x$, where $1 < x < 2$. Use devices from established manufacturers who have good quality assurance.
- Flicker noise (also called $1/f$, semiconductor, contact, or in resistors excess noise)-For this noise to occur, a DC current must be present. It is caused by different phenomena in different materials. Carbon resistors exhibit excess noise while metal film resistors do not. Use metal film resistors in low-noise designs.
- Johnson noise (also called thermal or Nyquist noise)—Resistors exhibit thermal noise caused by random thermal vibrations of charge carriers. Ideal reactive components do not suffer from thermal noise.
- Shot noise-Current passing through a semiconductor is the sum of many charge carriers bumping across barriers. This random process causes shot noise. The macroscopic current contains this white noise.

ciated with secondary amps, multiplexers, and so on often becomes negligible.

Of course, the noise on the first amplifier and front-end components is amplified with the information signal. Select an ultralow-noise amplifier for the front end.

Ultralow-noise amplifiers must be selected with care [3]. The amplifier's input voltage noise increases with decreasing frequency. Input voltage noise is amplified by the gain of the amp and shows up at the output.

The amplifier's input current noise is then converted to a voltage ($V_{noise} = I_{input\ noise} \times R_{source}$) and multiplied by the amp gain. If the sensor's source impedance is high, the amplifier's input current noise may become significant.

Thermal effects must be considered. Offset voltage drifts with temperature. A low-noise amplifier that dissipates a lot of power is apt to cause more headaches than an amplifier with more noise and less sensitivity to thermal conditions.

Internal noise may be minimized by selecting low-noise components. Metal-film resistors are lower noise than carbon types. Since "excess" noise in resistors is dependent on the current through the device, keep currents small. And thermal noise in resistors is dependent on the resistor value, so keep values small.

Although these two statements are paradoxical, getting the right tradeoff is your bread and butter [4].

Zener diodes exhibit avalanche noise and should be avoided. Modern band-gap references have lower intrinsic noise and can be employed in the same circuit topologies as zeners.

PCB LAYOUT

Just as bad PCB layout can destroy an excellent design, good layout can reduce unwanted parasitic effects. Never let critical low-noise PCB sections be designed by a CAD operator untrained in such layouts.

Tight component placement is crucial for minimizing EM1 susceptibility. The EDN article "Systematic Approach Makes Op-amp Circuits Resist Radiated Noise" explains how to determine maximum benefits from component placement [5].

Minimize trace resistance and inductance. Use big traces (60 mil) for all analog signals. Do not compromise. Find the space for 60-mil traces!

Keep traces on the front end of any differential amp close and symmetrical. This step minimizes thermal EMF effects and maintains the amplifier's CMR advantages.

A ground plane is extremely helpful. A dedicated ground plane provides shielding from EMI, maximizes trace capacitance, and provides a low-impedance current-return path. A multilayer PCB is an excellent method to implement a dedicated ground plane.

The cost of a four-layer PCB is about twice that of a two-layer PCB. However, the additional layers are worth their cost. One internal layer can provide the analog ground plane.

The additional signal layer may route power with huge traces (60-100 mil). It also helps with the real-estate requirements to route 60-mil analog traces. A four-layer PCB is highly recommended for low-noise circuits.

Do not route noisy digital traces near the low-noise section. If digital control signals are required, avoid parallel traces. Keep the digital signals on the opposite side of the ground plane from low-noise signals.

FR-4 [glass epoxy PCB material) is not a perfect electrical insulator, and without a guard ring around the amplifier inputs, some leakage current from adjacent traces will occur. See "Working with High Impedance Op Amps" [6] for details on this concept.

Thermal EMF can kill a good design. Thermal EMF results from a thermal gradient across a junction of dissimilar metals (a thermocouple). Each connection on the PCB is a parasitic thermocouple. Minimize connections. A copper/tin-lead junction has a 1–3-µV/°C thermoelectric potential. Creating an isothermal environment for the circuit eliminates the thermal EMF.

To facilitate an isothermal environment, put critical traces on the same side of the PCB. FR-4 is a reasonable thermal insulator. Traces on opposite sides of the PCB can be at slightly different temperatures.

Use components with a low-height profile (e.g., axial caps), so a 0.125" copper plate can be placed across all low-noise components. This thermal mass prevents abrupt temperature swings.

IC sockets should be eliminated from any low-noise design since they create unnecessary parasitic thermocouples. Connectors for sensor wires are often a necessary evil, but you should also eliminate them if possible.

Consider using SMT 0.1-µF caps to decouple the analog power supplies.

These tiny chip capacitors can be mounted on the solder side right under the analog ICs. This positioning maximizes their effectiveness.

Avoid SMT ICs. Offset voltage in analog components is strongly correlated to piezoelectric die stress. SMT packages don't have the physical stability that DIP (particularly CERDIP) packages offer.

CONCLUSION

Low-noise design is certainly challenging, perhaps somewhat arcane, but absolutely doable. Op-amps are still op-amps. Gain is still gain. And, simple is still better.

New components and test equipment have greatly eased the difficulties of low-noise design. However, attention to detail is still the engineer's responsibility.

Hopefully, the ideas presented here will assist you. Check out the references. Most of them are free. ▣

Bob Perrin designs nanovolt mixed-signal systems for Decagon Devices Inc., the world leader in high-speed precision water-activity instrumentation for food quality and safety assurance. He may be reached at nanobob@decagon.com.

REFERENCES

1. Keithley Instruments, *Low-Level Measurements Handbook,* 4th edition, 1993.
2. Crystal Semiconductor, *A/D Conversion ICs,* Vol. 1, Data Book DADBS2, 1992.
3. National Semiconductor, *Noise Specs Confusing!,* Application Note AN-104, 1990.
4. Peter J. Fish, "Intrinsic Noise," *Electronic Noise and Low-Noise Design,* McGraw-Hill, New York, NY, 72-90, 1994.
5. Bonnie Baker and Jerald Graeme, "Systematic Approach Makes Op-amp Circuits Resist Radiated Noise" EDN, 40 (15), 93, July 20, 1995.
6. National Semiconductor, "Working with High Impedance Op-amps," Application Note AN-241, 1991.

SOURCES

Analog Devices
One Technology Way
P.O. Box 9106
Norwood, MA 02062-9106
(617) 329-4700
Fax: (617) 461-4261

Burr Brown Corp.
P.O. Box 11400
Tucson, AZ 85734-1400
(502) 746-l 111
Fax: (520) 741-3895

Circuit Components, Inc.
2400 South Roosevelt St.
Tempe, AZ 85282
(602) 967-0624
Fax: (602) 967-9385

Design Tips App. Note
Cir-Quick, Inc.
12817 N.E. 126 Pl.
Kirkland, WA 98034-7721
(206) 823-7000
Fax: (206) 823-0381

Crystal Semiconductor Corp.
4210 S. Industrial Dr.
Austin, TX 78744
(512) 445-7222
Fax: (512) 445-7581

Keithley Instruments, Inc.
28775 Aurora Rd.
Cleveland, OH 44139
(216) 248-0400
Fax: (216) 248-6168

LeCroy Corp.
700 Chestnut Ridge Rd.
Chestnut Ridge, NY 10977
(914) 578-5985
Fax: (914) 578-6020

Maxim Integrated Products
120 San Gabriel Dr.
Sunnyvale, CA 94086
(408) 737-7600
Fax: (408) 737-7194

I R S

410 Very Useful
411 Moderately Useful
412 Not Useful

Light Probe for the Blind

FEATURE ARTICLE

Wayne Thompson

Light probes enable visually impaired people "see" light intensity via audio signals. All it takes to build this inexpensive device is a TSL235 converter, an AT-17 audio transducer, and a custom PC board.

For years, many people who are blind or visually impaired have used what are called *light probes*. This hand-held device produces an audio tone whose frequency corresponds to the intensity of the light it is pointed toward. The brighter the light, the higher the pitch of the tone.

Many light-probe designs have been implemented over the years with various features and degrees of usefulness. Most designs are small, battery powered, activated by push button, and include some sort of barrel for pointing toward a light source.

The good ones have a wide range of sensitivity and give instantaneous audible response to subtle changes in light intensity. They're used for detecting indicator lights on telephones, alarm panels, printers, modems, or any device in an office, factory, or home that uses lighted indicators to convey information to the user.

A popular vocation among persons who are blind is that of telephone receptionist. A light probe is often essential in determining which lines are busy, ringing, or on hold. The various flash patterns of the different lights on a multiline phone can be conveyed to a blind operator through corresponding audio tone shifts as each light is examined with the light probe. As lights flash on and off, tones sound high and low.

Similarly, radio disc jockeys and emergency dispatch operators-two other popular vocations of the blind—have a multitude of indicator lights they must detect.

Light probes also detect windows (during daylight hours) and can indicate whether overhead lights or table lamps are on. With some practice, you can even play tunes with a light probe by making slight yet precise movements in the probe's orientation between button presses. Right now, I'm working on "Twinkle, Twinkle, Little Star."

I HEAR THE LIGHT

Texas Instruments recently developed the TSL23x series of intelligent optosensors, as described in Tom Cantrell's article "IC de Light" (INK 62). The TSL235 light-to-frequency converter enabled us at the Kentucky Department for the Blind to redesign a light probe we've been building for consumers for over 15 years.

The new design is much smaller, lighter, and less expensive than all known designs. Its sensitivity and range are excellent, going from a 1- or 2-Hz clicking sound in near darkness to near ultrasonic in bright daylight.

The tiny size of the TSL235 three-lead package enabled us to create a very small device, not much larger than the two AAA batteries which power it. When activated, its current consumption is less than 2 mA, making for extremely long battery life under normal usage. The total cost for parts is only about $5, and a complete unit can be built in about 15 minutes (not counting the 24 hours of curing time for the barrel adhesive).

The circuit itself (see Figure 1) couldn't be much simpler, since all the work is done within the TSL235. The TTL-compatible output is sufficient to directly drive the small AT- 17 audio transducer from Projects Unlimited. Battery power and a switch complete the circuit.

There's a time-honored principle in rehabilitation engineering that goes, "Whenever possible, use a brick." Recent years have brought a wealth of high-tech solutions and products to the handicapped marketplace, including Braille printers, sophisticated wheelchairs, and talking computers. But, simple and cheap solutions to problems are still the first choice for rehab engineers.

The “use a brick” expression comes from the ultrasimple solution of wheelchair access to a table or desk— place a brick under each leg to raise the table height. Certainly, this project qualifies as an electronic brick, even though the technology behind the creation of the TSL235 is anything but simple.

As shown in Photo **1**, the TSL235, switch, and transducer are all mounted on a 2” x 0.9” printed circuit board. The board solders directly onto a AAA-battery holder which becomes, in effect, the finished case. (I said it was simple!)

Glue the barrel over the light sensor, pop in two AAA batteries, and drive your family crazy with squealing noises as you explore every light source in your home. Fluorescent lamps create a 60-Hz modulation buzz sound which is especially annoying!

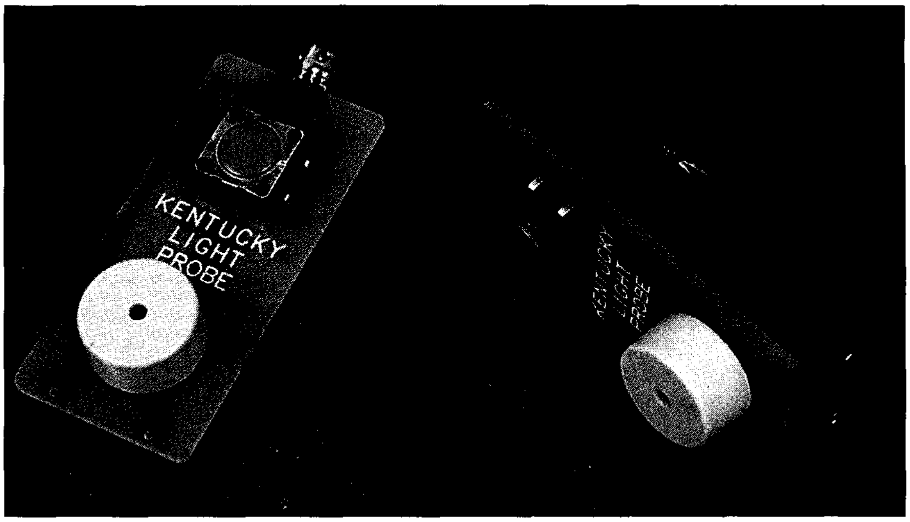

Photo 1—*The Kentucky Light Probe is a hand-he/d device that enables blind or visually impaired* **persons** *to perceive different light intensities by means of audio signals.*

CONSTRUCTION LITE

The Kentucky Light Probe is easily built. The parts list includes ordering information for each of its components.

First, cut the two plastic spacers off the bottom of the push button housing so that the switch lies flat against the PC board. Mount and solder the switch and audio transducer to the board. Cut off the protruding six component leads after soldering so they are nearly flush with the PC board.

Trim the three leads of the TSL235 so that only 1/4″ lead length remains. Position the TSL235 on the solder side of the board so the front sensing surface faces the component side and protrudes beyond the board edge. Line up the three leads with the three corresponding board traces and solder. The TSL235’s soldered position should be such that it can later be bent back 90” to meet the end surface of the battery holder, facing outward.

Next, trim the battery holder leads to 0.1”. Place the assembled PC board onto the battery holder, inserting the two power pins into the corresponding holes on the board. Then, swivel the board back **to** about a 45” angle, keeping the power pins in their holes. This move causes the battery holder’s pins to bend slightly.

Now, get your soldering-iron tip under the board for soldering the two power pins to the PC board pads. Insert the batteries and test the unit by pressing the push button. You may have to cover the sensor with your hand since the black vinyl barrel is not yet in place. With no barrel, normal room light sends the optosensor frequency into the ultrasonic range, fooling you into thinking it’s not working (although your dog may assure you that it is!)

Once you’re convinced that the soldered power connections are good, bend the PC board back down flat against the battery holder. Secure it with a spot of hot glue between the board and battery holder or temporarily use a rubber band to hold the two securely together. Avoid having battery-holder leads protrude through the PC board since the top of the PC board serves as an outer case surface.

Next, place a small dot of Loctite Black Max Gel adhesive on the back side of the TSL235 and bend it over 90” until it’s flat against the battery holder, facing out.

Finally, install the black vinyl cap used as a light barrel. This black vinyl material effectively blocks ambient light, and its flexibility ensures durability. A rugged barrel and adhesive are essential since light probes are carried in pockets, purses, and briefcases and are inevitably dropped.

Drill a %/‘-diameter hole in the end of the vinyl cap. The vinyl cap size (0.5” x 0.25″), coupled with the hole diameter of 5/64″, result in a highly directional barrel. These dimensions also calibrate the light-to-frequency ratio so the entire audio range is used for the light intensities normally encountered. The small hole enables the user to zero in on small indicator lamps mounted close together, such as those found on multiline telephones.

The best adhesive I’ve found for mounting the vinyl cap over the light sensor is a black cyanoacrylate gel such as Loctite Black Max Gel. Run a small bead of gel around the inside perimeter of the cap. Position and hold it firmly over the light sensor and flat against the battery holder for one minute or until it sets.

There should be no gaps between the vinyl cap and the battery holder surface which would let light leak into the sensor from the side. The black color of the adhesive helps achieve this “light-tight” seal.

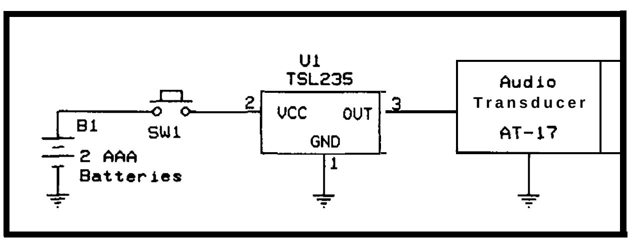

Figure 1—*The Light Probe incorporates the TSL235 light-to-frequency converter from Texas Instruments and an A T- 17 audio transducer from Projects Unlimited.*

Full curing takes about 24 hours. Ventilate the unit with a small fan to avoid the white-powder residue deposit common with this adhesive. Test the unit one more time before you leave the adhesive to cure.

NOW, GO BUILD ONE!

If potential sales justified it-there are about 500,000 blind persons in the U.S.-you could design a custom-molded plastic case. A battery-holder case looks a bit unfinished (although battery replacement couldn't be easier!). But, until that happens, maybe you can just do what we do: build the light probes one at a time for every blind person you know who needs one.

The PC board is available from the Kentucky Department for the Blind for $1. All other parts are commercially available. By the way, the blind community has long awaited an inexpensive color identifier. Could the TSL23x-series optosensors be used?

Warm up your soldering iron and call your local blindness rehabilitation agency for potential consumers in your area. Your low-tech skills could create a bright spot in the life of a blind person in your community. ❑

Wayne Thompson is an electrical engineer working for the past 14 years at the Kentucky Department for the Blind, a rehabilitation agency serving persons who are blind or visually impaired. He may be reached at Kentucky Department for the Blind, P.O. Box 757, 209 St. Clair St., Frankfort, KY 40601, by phone at (502) 564-4754, or by fax at (502) 564-2951.

CONTACT

Texas Instruments, Inc.
Literature Research Center
P.O. Box 172228
Denver, CO 802 17
(800) 477-8924, Ext. 3411
(303) 294-3747, Ext. 3411

I R S

413 Very Useful
414 Moderately Useful
415 Not Useful

PARTS LIST

TSL235 Light-to-Frequency Converter
Manufacturer: Texas Instruments
Vendor: Hamilton/Hallmark (800) 332-8638

AT- 17 Audio Transducer
Manufacturer: Projects Unlimited
Vendor: J.C. Hofstetter (513) 296-1010

SW412-ND Pushbutton Switch
Manufacturer: Omron B3F-4005
Vendor: Digi-Key (800) 344-4539

BH2AAA-PC-ND Battery Holder, 2-AAA, PC mount
Manufacturer: Memory Protection Devices
Vendor: Digi-Key (800) 344-4539

P265-ND Battery, carbon zinc, size AAA
Manufacturer: Panasonic UM4NX
Vendor: Digi-Key (800) 344-4539

Kentucky Light Probe Printed Circuit Board
Manufacturer: Southland Mfg. Co.
Vendor: Southland Mfg. Co. (606) 253-3066
(also KY Dept. for the Blind (502) 564-4754)

9753K35 Vinyl Cap, black stretchable
Vendor: McMaster-Carr (708) 833-0300

74765A84 Loctite Black Max Gel Adhesive
Vendor: McMaster-Carr (708) 833-0300

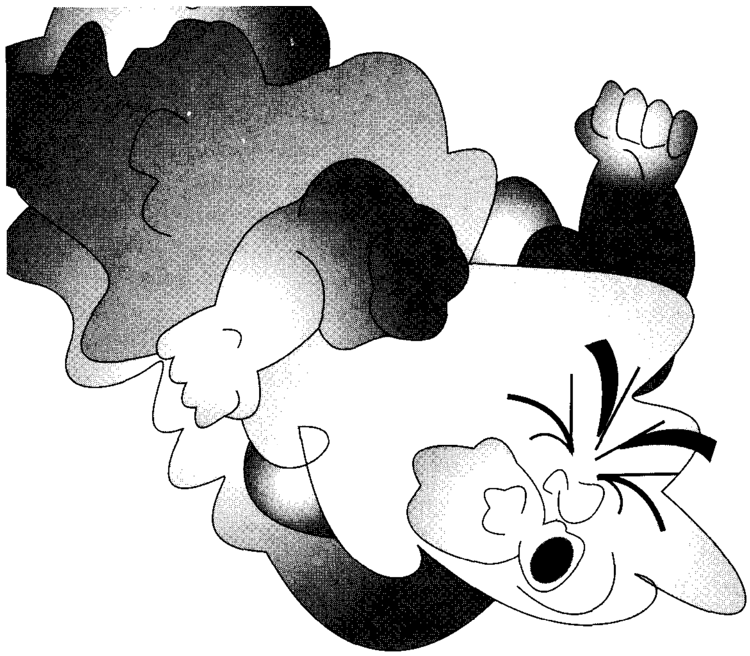

HA&BC

MARCH 1996

Innovations

IN HOME AUTOMATION & BUILDING CONTROL

edited by
Harv Weiner

HOME CONTROL SYSTEM

The HA1 **Omni Home Control and Security System** combines security and home control to provide a more convenient, secure, and energy-efficient living environment. The system controls security, temperature, lighting, and appliances and can be installed in new or existing homes.

Omni is interactive, so its security sensors can work double duty, providing the convenience of automatic lighting control and energy savings by setting back temperatures in unoccupied areas. Omni has modes of operation such as Day, Night, Away, and Vacation that determine the level of security required as well as the indoor temperatures and lighting levels. It's easy to select the desired mode using a control panel or Touch-Tone phone. Special modes can also be configured.

Working in conjunction with HAI's RC series of electronic thermostats, Omni can control temperatures within 0.5°F for optimum comfort and energy savings. By adjusting temperatures based on occupancy as well as time, the system provides greater comfort and efficiency compared to standard or programmable-setback thermostats.

Omni systems include telephone access in which any Touch-Tone phone (inside the home or away) becomes the control panel. The system gives a menu of choices in a clear human voice. The status of the system can be checked and settings of security, temperature, or lighting can be changed.

All functions of the system can be programmed on a timed basis (e.g., by time and day) or an event basis (door opened, motion detected, etc.) to provide automatic lighting and temperature control. Manual override is always available via local controls.

Home Automation, Inc.
2709 Ridgelake Dr.
Metairie, LA 70002
(504) 833-7256
Fax: (504) 833-7258

#510

PLUMBING CONTROL PANEL

Ultraflo Corporation has developed a solid-state control panel for its one-line, temperature-controlled, push-button plumbing systems. The panel not only includes new push-button switches, but it incorporates a new timing system which permits programming of shutoff, disposal, and hot-water safety options. The **Ultraflo system** eliminates faucet valves by using remotely operated low-voltage control to regulate water temperature and flow. Its single-line concept offers reduced maintenance and eliminates the traditional two-pipe system of hot and cold water.

The heart of the Ultraflo system is the central valve unit which is installed adjacent to the hot-water source (water heater or circulating lines). All water blending and flow control is accomplished within the valve unit and each outlet location. Connections between the valve unit and outlet locations are made with compression fittings, requiring only a wrench to install. Electrical connections use RJ-1 1-type phone jacks, and the switch can be installed in any convenient location.

The system also saves water and energy. The push-button, preset temperature saves water, time, and energy and provides complete safety from scalding water. The system uses only one ½″ or ⅜″ OD flexible line to supply each spout, which reduces the amount of water wasted when trying to adjust temperature.

Ultraflo Corp.
P.O. Box 2294 • 310 Industrial Ln. • Sandusky, OH 44870
(419) 626-8182 • Fax: (419) 626-8183

#511

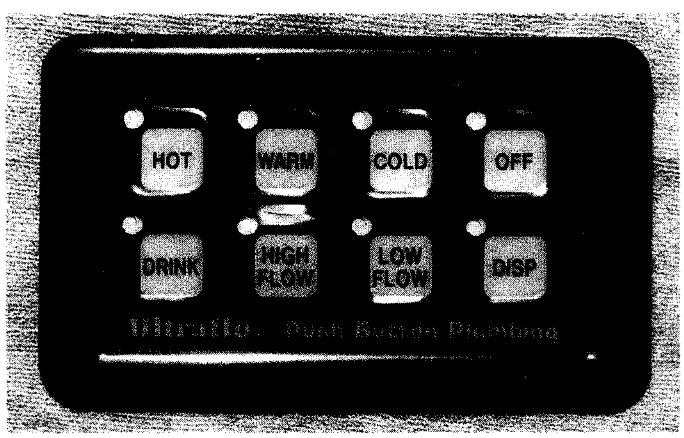

Inno Vations

ACCURATE ROOM-TEMPERATURE SENSOR

Automated Environmental Systems is shipping the **ATP-1000 Analog TempPlate,** a Decora-style temperature-sensing wall plate designed to provide accurate room-temperature monitoring in homes and office buildings. This low-voltage device is easily installed in a prewired, single-gang junction box or ring plate. The plastic body contains a low-profile metal probe which accurately senses local temperature. For accuracy, Template's temperature-sensing element is thermally isolated from the inner-wall temperature, which may be significantly different from the actual room temperature.

When provided with input power from +5 to +24 VDC, the Analog TempPlate outputs an analog signal proportional to temperature. The output signal is linear and scaled to produce +10 mV per degree Fahrenheit. The unit is capable of measuring temperature from 32°F to 212°F with an accuracy of ±1.0°F at 77°F. It is designed to drive a high-impedance load through hundreds of feet of shielded wire for true remote applications.

Applications for the Analog TempPlate include HVAC control (heating, cooling, damper control, and zoning), home automation devices (window and window-covering control), and data acquisition (temperature logging).

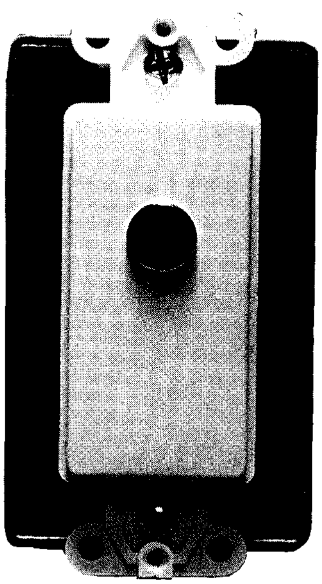

Automated Environmental Systems
4141 Ball Rd., Ste. 529
Cypress, CA 90630
(714) 826-7908
Fax: (714) 826-2908

X-10 TV CONTROLLER AND POWER-LINE IR

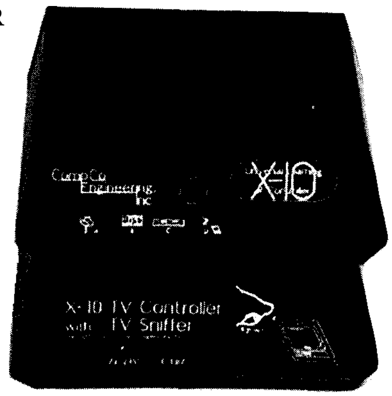

CompCo Engineering's **X-10 TV Controller** (XTVC) provides closed-loop communication with an X- 10 host controller. It monitors whether the TV is on or off and reports the state through the X-10 system to the host controller. No X-IO home-automation package is necessary to add this control to the TV.

The **TV Sniffer** adds volume and channel adjustments through the X- 10 control system and can also report manually or by IR remote to the host controller when the TV is turned on or off. When polled, the Sniffer reports the TV's status at any time. Using adjacent unit codes, the X- IO controller can change TV settings. The X-10 All Units Off command turns off the TV as well as lamps and appliances.

Remote speakers can also gain automatic control from the XTVC by setting an X-10 appliance module to the same code as the XTVC and plugging in the speakers. Whenever the TV is turned on or off, the speakers automatically are too.

A recent addition to the CompCo line is the **X-10 Universal IR Controller** (XUIRC), which combines the best features of X-l 0 and infrared (IR). IR is popular in home automation because a wide range of electronics are ready to be controlled with IR. The XUIRC "learns" IR from remotes and responds to X- 10 by sending almost any IR command, which places off-the-shelf IR components under X-10 control.

CompCo X- 10 controllers eliminate the need to run separate wiring for IR transmitters in other rooms. The XUIRC includes the same closed-loop communication features as the XTVC, and also includes a TV Sniffer probe.

The XTVC with TV Sniffer sells for $149. The XUIRC retails for $299.

CompCo Engineering, Inc.
611 Parkway, Ste. D4
Gatlinburg, TN 37738
(423) 436-5189
Fax: (423) 436-5189

#513

A Wind Speed and Direction Interface for the HCS II

JOHN MORLEY

In addition to describing a wind speed and direction interface for the HCS II home control system, John presents background information on the various techniques for measuring wind speed and direction. And, unlike the golfer, his finger stays dry and there's no grass in his hair.

If you're like most people, many of your leisure activities depend on the weather. For outdoor activities such as golf, sailing, and biking, good weather is essential.

As an avid private pilot, this situation is especially true for me. Sure, I can fly when the weather is not good, but I much prefer to fly when the sky is blue and the sun is shining. These conditions have become so important to me that I sometimes find myself checking the weekend weather forecast as early as Tuesday!

As a result of my personal "dependence" on good weather, I've developed a keen interest in collecting my own weather data, making my own weather observations, and forecasting the weather. For many years, I've gathered most of my weather data via traditional outlets such as NOAA weather radio and local TV and radio broadcasts.

Over the past few years, easy access to the Internet has enabled me to supplement this data with instant access to both national and regional weather data and forecasts. Even more recently, I've added weather-data-measurement capability to my HCS II system. I can now directly collect and analyze my own weather data.

Until now, my weather data-collection system consisted of sensors for measuring outdoor temperature, relative humidity, and barometric pressure. My most recent additions—transducers and signal conditioners—measure the speed and direction of the wind.

This article details the design, construction, and calibration of the wind-measurement instruments.

THE BASICS

Wind is a phenomenon caused by large moving masses of air molecules. These molecules comprise the gaseous atmosphere that surrounds our planet.

Heat is the driving force behind the movement of air molecules in our atmosphere. Solar radiation heats the surface of the earth, which in turn heats the air closest to it.

This heated air begins to rise, forming convective air currents (i.e., vertical columns of rising warm air). As the heated air rises, it gradually loses its heat, falls away, and is replaced by the continually rising stream of warm air from the surface.

Global circulation results when the intense sunlight striking the equator causes this convective upwelling of heated air molecules. These air molecules flow north and south toward the poles, cooling and sinking en route. Cooler air from the polar regions is in turn drawn into the solar-powered furnace at the equator.

This natural flow creates a gradient of air pressure, with low-pressure areas near the equator increasing to high-pressure areas near the poles. Finally, the effects of the earth's rotation about its axis work to ensure that the effects of atmospheric circulation, and varied atmospheric pressure, are distributed over a wide geographical area.

These factors contribute in large part to the continually changing weather patterns we experience every day. Moving air masses (wind) are most often quantified in terms of their relative direction and velocity. Because the type of weather we experience is closely linked to these moving air masses, a great deal of meteorological information can be gleaned from wind-speed and wind-direction measurements.

For instance, the passage of a weather front (cold or warm) is always followed by a change in the wind direction, temperature, and barometric pressure. Wind speed can also be an excellent indicator of current and future weather conditions.

In general, constant wind direction and low to moderate (not gusting) wind velocity indicate a stable air mass and thus fair weather. High wind velocity, gusting, and variable wind direction indicate an unstable air mass and unsettled weather.

Your own observations and the weather-forecasting experience you acquire will likely be the most powerful tools you have to accurately gauge the implications of the weather data you observe. The relationship between this data and the actual weather is left to you for further investigation.

If you thought that adding weather measurement sensors to your HCS II system was beyond you, think again! I'll show you the steps I took to design and build a complete wind speed and direction measurement interface for my HCS II system.

So, gather the required parts and warm up your soldering iron, while I show you what's necessary to add this interface to your own home control system.

JUST HOW DO I MEASURE?

The two instruments commonly used to measure wind speed and wind direction are the anemometer (wind speed) and the wind vane (wind direction). The wind vane, illustrated in Figure 1, generally consists of a weighted pointer connected to a small flat plate. This combination is attached to a rotating vertical shaft.

The force of air molecules striking the flat plate causes the shaft to rotate, aligning the plate with the direction of the local air flow. The weighted pointer serves as a damper to reduce the effects of small variations in actual wind direction, as well as providing a visual indication of the direction the instrument is pointing for alignment purposes. A direct wind-direction readout (or a potentiometer for remote readings) is usually attached to the rotating shaft of the wind vane.

For meteorological purposes, wind direction is universally expressed in terms of compass heading degrees (O-360"). By the way, wind direction is always expressed in terms of the direction from which the wind is coming, not the direction in which the wind is going! The term *anemometer* is derived from the Greek word *anemo*, meaning *wind*, and refers to any instrument used to measure wind velocity.

There are several common units of measure for wind speed. Meteorologists in the U.S. generally express wind speed in terms of miles per hour (MPH), based on a statute mile. Certain disciplines, particularly marine and aviation, express wind speed in terms of knots (or nautical miles per hour). One kn is equal to 1.15 MPH.

In countries where the metric system is used, wind speed is often expressed in terms of kilometers per hour. One km/h is equal to 0.625 MPH.

A few popular anemometer types are reviewed below.

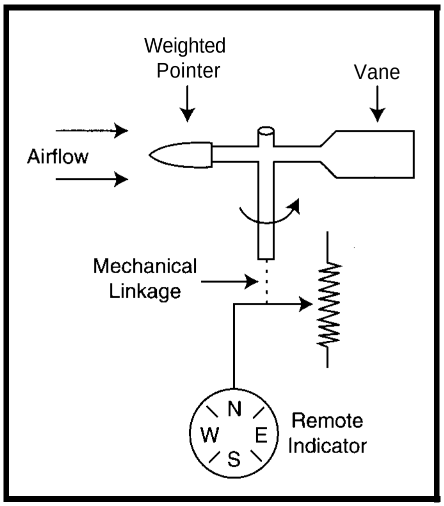

Figure 1: *The wind vane aligns itself with the local air stream to indicate wind direction.*

PITOT ANEMOMETER

Figure 2a illustrates the pitot anemometer, one of the simplest instruments that can be used to measure wind speed. This anemometer consists of a simple pitot tube and a pressure-measurement system. It is based on the principle that moving air molecules flowing into the pitot orifice exert a dynamic pressure proportional to the velocity of the air molecules in the tube.

To determine wind velocity, the pressure of the airstream is measured relative to the pressure of static air. This measurement is often made using a mechanical diaphragm (thin elastic partition) and indicator system.

As air velocity increases, air molecules strike the diaphragm, exerting an increasing amount of pressure and causing it to expand. The amount of displacement depends on the amount of pitot pressure and the prevailing atmospheric (static) pressure.

The diaphragm is linked via levers and gears to a pointer system which directly displays the air velocity. This measurement system is calibrated and the display indicator is marked in the appropriate units using a known wind-speed standard.

Recently, it has become more common to measure this pressure using a solid-state pressure sensor. This type of anemometer has the disadvantage of being very directional. Components of air velocity not in line with the axis of the pitot measurement tube are measured inaccurately, if at all.

For this reason, the pitot anemometer is usually found in applications in which the direction of the air velocity is fixed. Applications include air-velocity measurement in wind tunnels and air-speed indicators in airplanes.

HOT-WIRE ANEMOMETER

For very accurate laboratory-type measurements, the hot-wire anemometer (Figure 2b) is often used. This anemometer consists of an electrically heated wire on the end of a sensor probe and an electronic control module.

Air velocity is determined by measuring the amount of electric current it takes to maintain a preset constant temperature in the heated wire as air molecules pass over it.

As air molecules pass over the heated wire, they tend to cool the wire (forced-convection cooling). The hot-wire anemometer's control circuitry responds to this cooling (generally measured with solid-state sensors) by increasing the heating current in the wire. As air velocity increases, the heating current in the wire also increases by an amount proportional to the velocity.

Air velocity can then be determined by measuring the difference in the magnitude of the control signal (over zero airflow conditions) used to determine the heating current in the hot-wire. This magnitude is then applied to a calibration curve that is temperature compensated to account for the effects of ambient temperature changes.

The hot-wire anemometer is one of the most accurate means of measuring air velocity. It is also quite expensive,

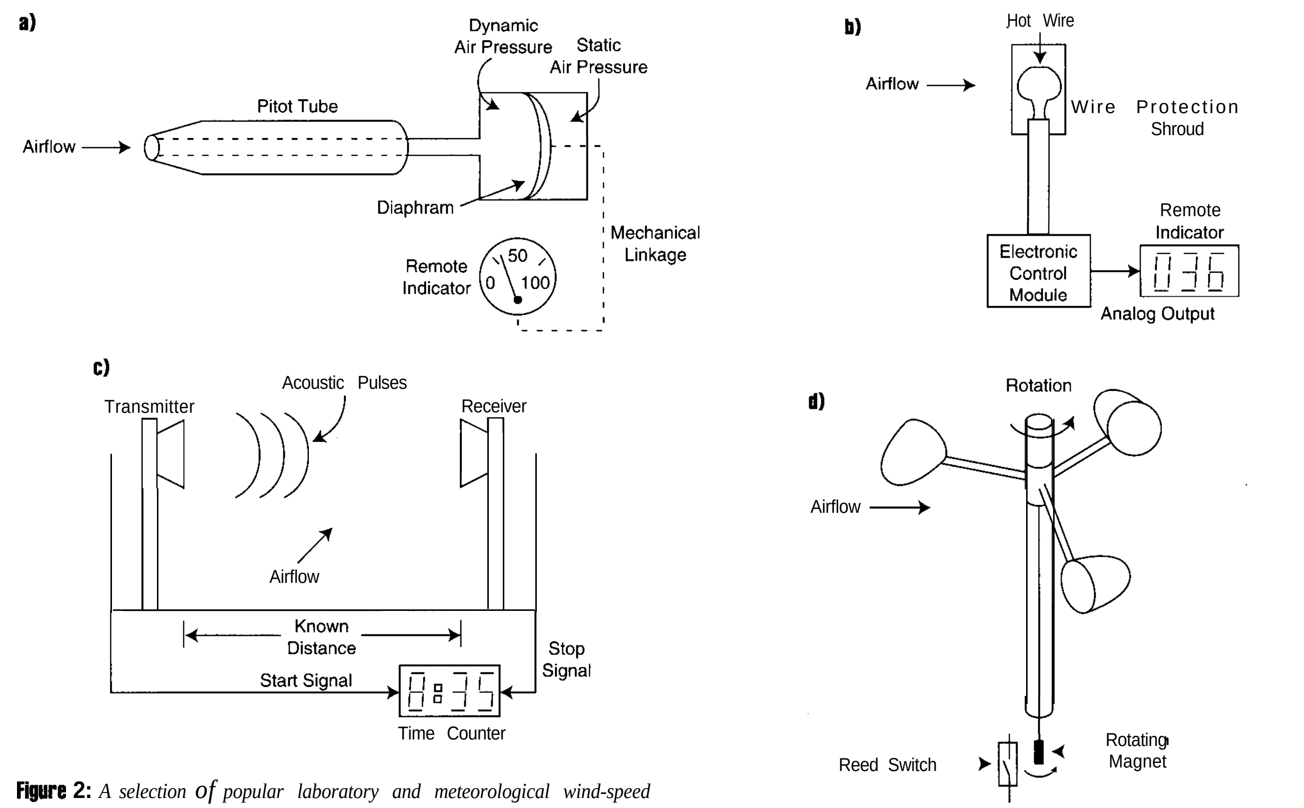

Figure 2: *A selection of popular laboratory and meteorological wind-speed transducers includes: (a) pitot anemometer, (b) hot-wire anemometer, (c) acoustic anemometer, and (d) rotational anemometer.*

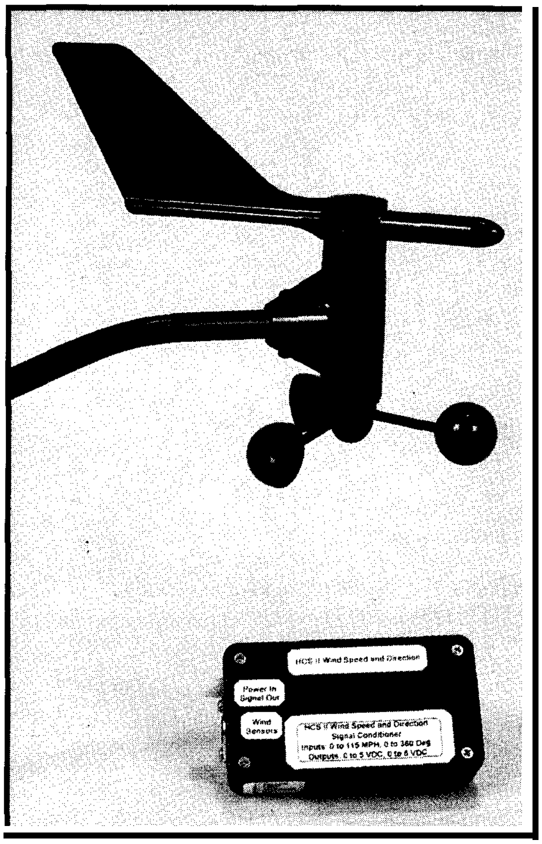

Photo 1: *These small, rugged wind transducers are conveniently mounted on a common offset bracket for ease of installation.*

which limits its application to the laboratory environment.

ACOUSTIC ANEMOMETER

The acoustic anemometer works by exploiting a very well-known principle: moving air molecules affect the speed of an acoustic sound wave. This type of anemometer (Figure 2c) usually consists of a sound-energy source (transmitter), and a sound-energy collector (receiver), separated by a precisely measured distance of several hundred feet.

Sound energy is emitted from the transmitter and shortly thereafter is received at the collector. The time required for the sound signal to reach its destination is proportional (among other things) to the velocity of air molecules moving in the space between the transmitter and receiver elements.

Due to its widely spaced construction, an acoustic anemometer can capture a more general picture of wind activity than other anemometer types. Unfortunately, the speed of sound is also affected by changes in air density due to temperature and barometric-pressure changes. Again, due to its construction, this type of anemometer is much more likely to be affected by these changes and must be compensated to obtain accurate results.

Occasionally, acoustic anemometers use a rotating energy source and multiple sound receivers. The wind-profile information that this instrument provides makes it well suited to applications such as the detection of violent wind shear near airports.

ROTATIONAL ANEMOMETER

The rotational anemometer (Figure 2d) should be familiar to you, as it is the most common type used for meteorological measurements. This anemometer consists of a windmill, propeller, or as is most often the case, three semiconical cups attached to a rotating horizontal shaft.

Moving air molecules striking this anemometer exert a force on the cups, causing the shaft to rotate about its axis. As air velocity increases, the anemometer shaft's rotational velocity increases proportionately. The shaft is often directly coupled to an electric generator which measures shaft rotational speed and, thus, wind speed.

In most generator-type rotating anemometers, the magnitude of the generator's AC voltage output increases with shaft frequency. A precision rectifier and filter can be used to convert the AC voltage generated by the rotating shaft into a more useful DC voltage. This DC voltage is then directly displayed on a voltmeter calibrated to wind speed.

More common today, however, is a mechanical or electronic switch output. In these anemometers, there is at least one (and often more) switch closure(s) for each rotation of the anemometer shaft. The frequency at which switch closures occur is proportional to the wind speed. This frequency can be converted to an analog voltage for measurement, or a microcontroller or computer can measure wind speed directly.

HCS II INTERFACE

The HCS II wind speed and direction interface is based on a commercially manufactured anemometer and wind-vane transducer assembly like the one in Photo 1. These instruments were chosen primarily for their robust construction, which promises many years of maintenance-free operation. They also came the closest to satisfying my design goals for availability, performance, and price.

As you can see in the photo, the two instruments are paired on a bracket designed to be offset-mounted to a vertical mast. The anemometer and wind-vane transducer assembly is supplied complete with mounting hardware and 100' of cable.

The wind-vane shaft is internally coupled to a 20-k potentiometer for measuring shaft position and wind direction, This potentiometer has a complete 360" rotation with no rotational stops. The anemometer shaft is internally coupled to a magnetic-reed switch. Each revolution of the anemometer produces one contact closure of the reed switch.

The HCS II Supervisory Controller has a number of S-bit analog (ADC) inputs designed to be connected to analog-type measurement sensors.

Other analog inputs can be added if you use network modules. These analog inputs are designed to be interfaced to ground-referenced 0–5-VDC analog signals. The HCS II, however, is not capable of measuring the frequency of a periodic input signal (digital) with the accuracy required for this interface.

Clearly, our design challenge is to implement the circuitry required to interface the anemometer and wind-vane transducer outputs to the HCS II analog measurement inputs. In the case of wind direction, it's actually quite easy.

To be compatible with an HCS II analog input, a signal is needed that varies from 0 to 5 VDC and is proportional to wind direction. Remember, the wind-vane shaft is directly connected to the wiper of a 20-k potentiometer. Therefore, the position of the pot wiper always indicates wind direction.

It's a simple matter to connect the ends of the potentiometer to ground and to +5 VDC. The wiper voltage then varies from 0 V (000") to 5 V (360") as the wind vane is rotated. The relationship between the wiper output voltage and wind direction can be determined by:

$$\text{WindDirection} = \frac{V_{out}}{5\text{ V}} \times 360"$$

This relationship assumes that the potentiometer is completely linear. According to the manufacturer, the linearity tolerance for the potentiometer is ±5% or less.

The interface circuitry for measuring wind speed is somewhat more complicated. Remember that the anemometer shaft is coupled to a switch. This switch closes once for every revolution of the anemometer.

The first step in the design of the wind-speed interface is to calibrate the anemometer to correlate wind speed versus switch-closure frequency. To accomplish this step, I mounted the wind transducers to the spare-tire carrier at the rear of my truck. Then, I drove the vehicle on a virtually windless morning at speeds spanning the input range of the anemometer while measuring the switch-closure frequency.

Before attempting this experiment, I also had my vehicle's speedometer professionally calibrated since my new tire diameter differed from the original tires. The recalibration resulted in a fixed indicated speed offset of about +3 MPH.

The results of my data collection are presented in Figure 3. The manufacturer subsequently verified my findings-one switch closure per second (or 1 Hz) equaled 2.33 MPH. Thus, over the wind-speed range of interest (O-l 15 MPH), the switch closure frequency varies from about 1 Hz (-2.3 MPH) to 50 Hz (- 115 MPH).

There are many ways to convert a variable frequency signal into a DC voltage. This type of circuit, often known as a tachometer circuit (due to its common use with engine tachometers) is implemented most often using a single-chip FN converter integrated circuit. The circuit consists of an input comparator, one-shot, and integrator.

Each time the input signal crosses the comparator threshold, it activates the one-shot, which switches a precise amount of charge into the integrator input for a measured time period. As the input signal

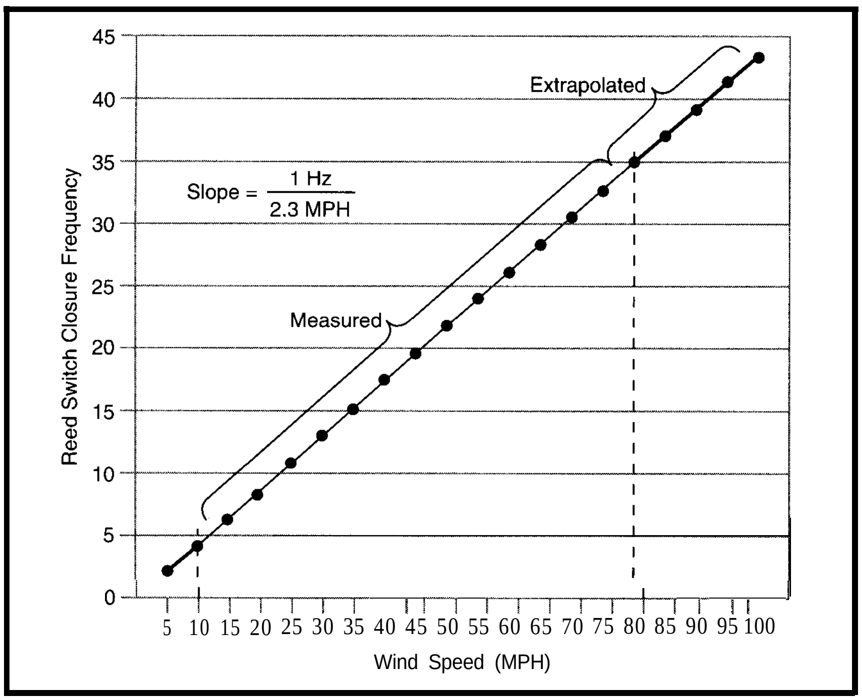

Figure 3: *A linear relationship exists between reed-switch closure frequency and wind speed. All data was measured experimentally (see text for details).*

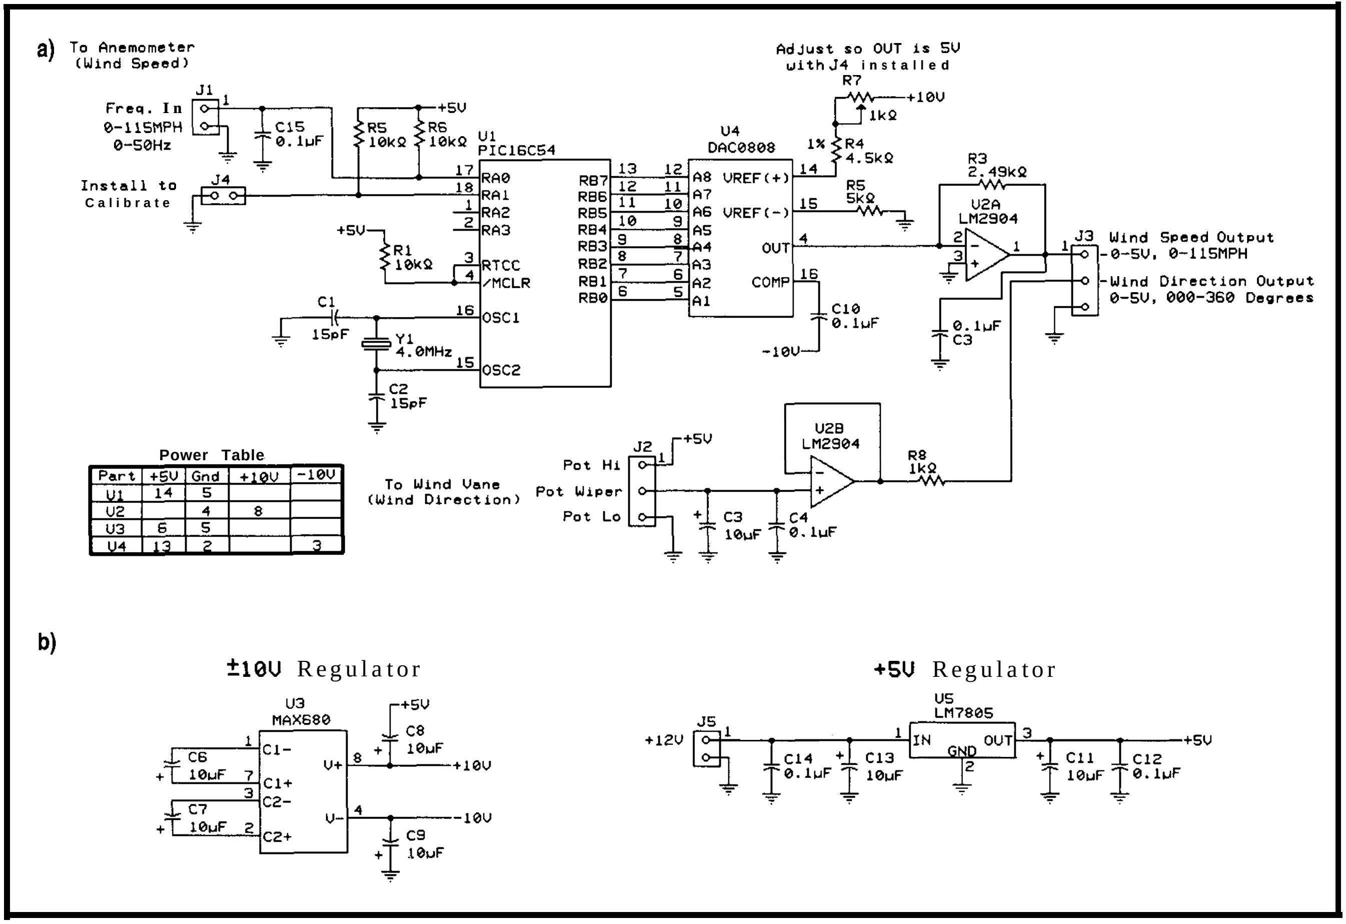

Figure 4: *The wind-speed and direction sensor outputs are converted to analog voltages (O-5 V) to be measured by the HCS II.*

increases in frequency, the charge injected into the integration capacitor increases proportionally. The result is an average output voltage from the integrator proportional to the input-signal frequency.

Unfortunately, a number of minimum input-signal characteristics must be observed for a circuit of this type to operate correctly. For this application, the most important specification is the input frequency range.

Most FN converters generally operate best with frequency inputs greater than 10 kHz. Although they continue to operate at much lower frequencies, the amount of the input frequency that appears at the output (carry-through frequency) goes up dramatically as frequencies decrease.

You can make allowances for reduced input-frequency operation by varying the amount of charge injected by the one-shot into the integrator or by enlarging the integration capacitor. However, these changes can lead to unacceptable circuit-response time and circuit instability.

From Figure 3, it's clear that the frequency range of the switch output, across the anemometer's wind-speed range, is substantially less than acceptable. Despite this problem, I decided to try a few experiments with a FN converter that I had on hand (Analog Devices AD650) just to see how bad the situation would be.

As I expected, the output-voltage ripple at the lowest expected input frequency was extremely high, approaching almost 1 Vp-p! Clearly, I needed to measure wind speed (switch-closure frequency) using an entirely different technique, since the frequencies being measured were extremely low.

Obviously, this job was better suited to a digital counter or microcontroller! I chose Microchip's PIC16C54 due to the circuit design's simplicity, low parts count, and program flexibility.

In the approach I took, the wind-speed transducer is connected directly to an input bit of the PIC microcontroller. An 8-bit D/A converter configured for an output range of O-5 VDC is connected to eight of the PIC's general-purpose I/O bits, which are configured as outputs.

To measure wind speed, the PIC simply counts the number of

Module Interface Connectors	
Wind Instrument Interface Connector (5-Pin DIN Female)	HCS II Interface Connector (4-Pin Molex Female)
Pin 1: +5 V	Pin 1: +12 V
Pin 2: Wind Direction	Pin 2: Gnd
Pin 3: Gnd	Pin 3: Wind Direction
Pin 4: Wind Speed	Pin 4: Wind Speed
Pin 5: NC	

Table 1: The wind speed and direction interface module uses standard connectors to make disconnecting the unit simple.

Listing 1: *XPRESS code can be used to display the wind speed and direction on an LCD-Link and in the HOST console window.*

```
DEFINE Ch5RawData=Variable(5)
DEFINE Ch5Volts=Variable(6)
DEFINE WindSpeed=Variable(7)
DEFINE Ch6RawData=Variable(8)
DEFINE Ch6Volts=Variable(9)
DEFINE WindDirection=Variable(10)

BEGIN

IF Reset THEN
   Timer(60) = ON
END

IF Timer(60)>=1 THEN
! Here I read analog input channels 5 and 6
! and convert to standard wind speed and direction units.
   Ch5RawData = ADC(5)
   Ch5Volts = (Ch5RawData * 50) / 256
   WindSpeed = (Ch5Volts * 23) / 10
! Wind speed is 0 MPH at 0 V, 100 MPH at 5 V.
   Ch6RawData = ADC(6)
   Ch6Volts = (Ch6RawData * 50) / 256
   WindDirection = (Ch6Volts * 72) / 10
    IF WindDirection = 360 THEN
       WindDirection = 0
    END
! Wind direction is 000° at 0 V, 360° at 5 V.

! Here I display wind speed and direction
! on the LCD-Link and in the V3.0 HOST message window.
   LCD(0)  = "Wind Spd: %PO MPH\n", WindSpeed
   Console = "Wind Spd: %PO MPH\n", WindSpeed
   LCD(0)  = "Wind Dir: %PO deg.\n", WindDirection
   Console = "Wind Dir: %PO deg.\n", WindDirection
   Timer(60) = ON
END
```

anemometer switch closures during a predetermined time period (frequency). This count is then scaled to obtain the 8-bit digital value which drives the DAC. The scale factor ensures that the digital output is equal to -255 when the input frequency is equal to 50 Hz (-115 MPH).

The DAC's output voltage is always proportional to the wind speed as measured by the anemometer. The software is programmed to measure frequencies up to 50 Hz in 1-Hz increments. Thus, wind speeds up to 115 MPH can be measured in increments of 2.3 MPH. Therefore, at a wind speed of 0 MPH, the output voltage of the circuit is 0 V, and it is -5 V with a wind speed of ~ 115 MPH.

The relationship between the DAC output voltage and wind speed can be found by:

$$\text{WindSpeed} = \frac{V_{out}}{5\text{ V}} \times 115\text{ MPH}$$

Using a microcontroller, you can measure switch-closure frequency (and thus, wind speed) with greater resolution than 1 Hz. In my application, however, this level of precision was unwarranted.

Besides, the nature of wind-speed measurements is that they tend to fluctuate at random intervals. To compensate, many wind-speed reporting devices average the speed data recorded over many seconds. Future versions of the software may include a user-programmable averaging time period.

CIRCUIT DESCRIPTION

The circuitry required to interface the wind speed and direction transducers is shown in Figure 4a.

As I mentioned earlier, the easiest method I've found for generating a O-5-V signal with the wind-vane transducer is to connect the ends of the 20-kΩ potentiometer in the sensor directly to GND and +5 V. The potentiometer's wiper voltage therefore varies between 0 and 5 V as the wind vane moves from 000" to 360".

The wiper voltage is connected to op-amp U2b, which is configured as a voltage follower. This amplifier provides impedance buffering for the potentiometer and a constant source impedance for the HCS II A/D converter.

Capacitors C3 and C4 filter the potentiometer voltage and reduce rapid voltage fluctuations. Should the potentiometer become disconnected, resistor R8 limits the output current that the amplifier would source into the overvoltage-protection diodes of the ADC on the HCS II.

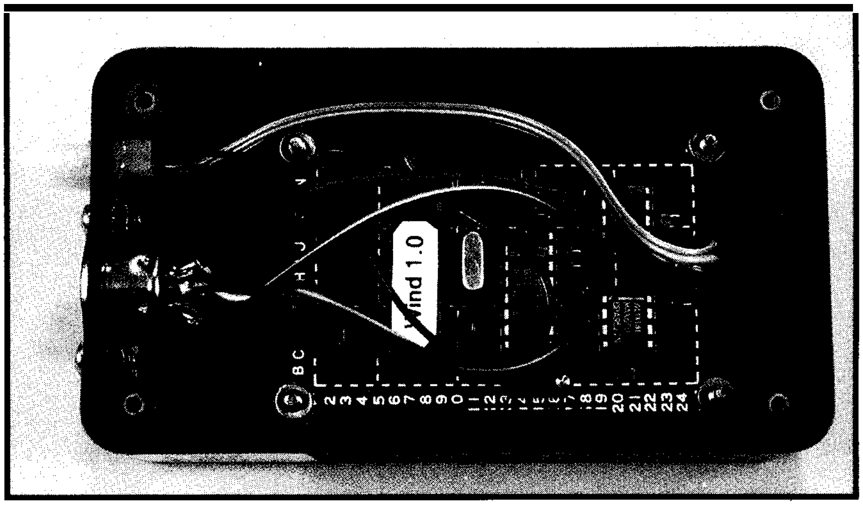

Photo 2: *The completed transducer interface prototype is intended to be mounted indoors, protected from the elements*

WIND SPEED

The anemometer switch input connects directly to the microcontroller's (U1) RAO I/O bit, which is configured in software as an input. The input is pulled up with R2, and C 15 provides a measure of debounce protection (τ = 0.0001 s).

The other end of the anemometer switch is connected to ground. The microcontroller's RA 1 I/O bit, also configured as an input, is connected to a configuration jumper. Connecting this jumper grounds the input and puts the microcontroller into calibration mode.

The PIC crystal (Y 1) is a 4.00-MHz unit, selected to provide a 1 -µs instruction-execution speed in software. This value was selected to help make precise time calculations more convenient.

An 8-bit DAC is connected to eight of the microcontroller's general-purpose I/O bits (RB0–RB7), which are configured as outputs. The DAC's current output is converted to a voltage by op-amp U2a. The converter and op-amp are configured so that the amplifier's output is 0 V (0 MPH) when the DAC digital input is OOH, and is 5 V (115 MPH) when the input is FFH.

Potentiometer R7 adjusts the full-scale (5-V) output of the DAC circuit. Voltage regulator U5 (LM78L05) provides the circuit with a regulated source of +5 V while U3 (MAX680) is a charge-pump voltage converter used to supply ±1 0 V to the DAC and op-amp, as shown in Figure 4b.

The recommended connector pinout information for the wind speed and direction interface module is depicted in Table 1.

INSTALLATION AND CALIBRATION

Installing and calibrating the HCS II wind speed and direction interface circuit (see Photo 2) is relatively straightforward. The anemometer and wind vane should be permanently mounted to a solid structure of your house, as close to the roof as possible.

In my installation, the sensors are attached to a 5' steel mast, which is mounted to the waste-vent pipe on the roof using a mounting bracket from a local electronic parts distributor. It is also possible to find brackets that can be bolted to the side of the house or attached to your chimney.

Remember, it's generally a good idea to locate wind sensors as far above the roof and as far away from other structures (ridge lines, antennas, etc.) as possible. That way, you minimize wind blockage, turbulence, and other effects of these obstructions.

Next, you need an accurate compass to align the wind vane. I found it helpful to turn the wind vane to a known position such as south (180°) during the alignment procedure. Then, temporarily fix the wind vane in the known position with a piece of tape.

Now, using the compass, determine the direction to the preset heading and rotate the complete sensor assembly until the wind-vane pointer is facing the proper direction. Tighten the mounting hardware to ensure that the wind transducers won't move after alignment.

For the wind-speed interface circuitry, the only calibration needed is to adjust the DAC full-scale output. To make this adjustment, connect jumper J4 on the circuit board and adjust R7 so that pin 1 of the op-amp U2 is 5.0 V. The wind-speed interface circuit has no other adjustments.

The manufacturer provides the calibration information (1 Hz = 2.3 MPH), and you can independently verify it. This conversion constant is hard coded into the microcontroller firmware.

Being an eternal skeptic, however, I borrowed a very expensive wind-speed indicator during initial installation of my wind-speed interface to check the output of my circuit. I was pleasantly surprised to find that at any given wind velocity, the difference between the calibration standard and my circuit didn't vary by more than a few percent.

XPRESS CODE

The example code in Listing 1 details the XPRESS programming language statements necessary to display the data collected using the wind speed and direction interface. Every 60 s, the XPRESS code reads the wind speed and direction by measuring the voltages connected to the HCS II channels 5 and 6 analog inputs.

These voltages are then converted to standard wind speed and direction units. The way this process is accomplished is somewhat unique. Because the XPRESS language doesn't support floating-point mathematical operations, all calculations are made using "scaled" integers.

By carefully controlling the placement of the displayed decimal point, the proper wind speed and direction values are represented and displayed correctly at the LCD-Link or console message window. By dealing with scaled integers, we retain the most-significant portion of our data otherwise lost to round-off after performing integer math operations.

In this example, the raw ADC counts (O-255) of channels 5 and 6 are first converted to volts. The apparent range of computed voltages is O-50 V, as a result of scaling. Computed voltages are then converted to standard wind speed and direction units.

For wind speed, we simply multiply the computed input voltage range by 2. This calculation gives a wind-speed range of 0 (lower limit) at O-V input to 100 (upper limit) at 5-V input.

For wind direction, multiply the computed input-range voltage by 72 and divide by 10. This calculation gives a wind direction range of 0 (lower limit) at O-V input to 360 (upper limit) at 5-V input. In meteorological terms, wind direction is most often represented as ranging from 000" to 359" (000° = 360"). So, when the computed wind direction is 360°, reset the value to 000".

TAIL WIND

As with almost all data collection, interpreting results can be harder than collecting them. Weather data in particular is often just like that. In fact, it's been said many times that weather forecasting is the only profession that allows you to be wrong more than 50% of the time and still keep your job!

Weather forecasting seems to be much more an art than a science. You'll find, however, that the more data you collect and the more observations you make, the better your ability to recognize recurrent weather patterns.

Weather "signatures" begin to emerge as you identify the characteristic conditions for the weather you currently have or will have in the future. This is not to say that you won't occasionally be deluged during the summer barbecue you promised would be beautiful. My suggestion-play it safe. Keep an umbrella handy while you're still wrong 5 1% of the time!

John Morley is the senior electrical engineer on the staff of a small Boston area manufacturer of custom electronic test equipment. His primary responsibility is the design of instrumentation used to measure the thermal properties of packaged semiconductor devices, and the reliability of electrical interconnects. John may be contacted at endeavor@ usa1.com.

SOURCES

ADC0808 A/D converter, D/A converter, operational amplifier
National Semiconductor Corp.
2900 Semiconductor Dr.
P.O. Box 58090
Santa Clara, CA 95052
(817) 468-6300
Fax: (8 17) 468-6935

Anemometer and wind-vane sending unit
Amateur Electronic Supply
5710 W. Good Hope Rd.
Milwaukee, WI 53223
(414) 358-0333
Fax: (414) 358-3337

Charge-pump voltage converter and MAX680
Maxim Integrated Products
120 San Gabriel Dr.
Sunnyvale, CA 94086
(408) 737-7600
Fax: (408) 737-7 194

HCS II
Circuit Cellar, Inc.
4 Park St.
Vernon, CT 06066
(860) 875-275 1
Fax: (860) 872-2204

PIC16C54 microcontroller
Microchip Technology, Inc.
2355 W. Chandler Blvd.
Chandler, AZ 85224-6199
(602) 786-7200
Fax: (602) 899-9210

I R S

416 Very Useful
417 Moderately **Useful**
416 Not Useful

Designing and Evaluating the User Interface

Make your Design User Seductive

ORRIN CHARM

Although average consumers can handle automobiles and **ATMs,** they're stumped by VCRs, security alarms, or office phones. Orrin traces these difficulties to the user interface. He shows how to minimize such dilemmas for interfaces like remotes, video monitors, telephones, control panels, and voice-control products.

The average consumer has no difficulty driving an automobile or getting cash from an ATM, yet is stumped by a VCR, security alarm, and office phone system!

Why?

Frequently, it's the user interface.

Before looking at how to minimize such dilemmas in product design and selection, I review the user interfaces available on the market today and how to determine which are appropriate and for what kinds of applications.

DESIGNING FOR HOME

In the past few years, we've seen many new products, services and concepts bring the power of computers and information technology to the average home. Although we accept this as good and inevitable, as masters of technology, we tend to be comfortable with computer interfaces and the learning curves required to use them.

For the average homeowner, however, such devices are threatening. Offered a choice between learning an unfamiliar technology or bypassing it (e.g., going to the wall to flip the light switch), most homeowners inevitably choose the familiar.

It's slightly different in the business world. There, computers have gained a foothold mainly by force. Anyone entering the workforce in a clerical or office-support environment has no choice. They must learn standard software.

Only upper management has been able to resist computers. However, as they find themselves shut out from the information around them, even they are taking the big step.

At home. though, most people can choose to accept or reject technology. When X-10 devices first appeared, I installed them at home. I was thrilled, but my roommate threatened to throw them and me out if she couldn't turn on the lights the "normal" way!

We accelerate acceptance of new methods by forcing people to use them. If you could convince people to buy homes without conventional switches, they would soon use the new devices. In fact, it seems to me that home automation gains a foothold just this way: sell its benefits to the home builder.

Eventually, the products that are successful seduce potential owners with easy accessibility and a promise of rewards. The development of user-seductive human-to-machine interfaces is undoubtedly the most critical milestone in creating a viable market for home systems.

WHAT IS A USER INTERFACE?

Although initially defined by Webster as a common boundary between two bodies, spaces, or phases, its secondary definition fits our purposes better. In this context, it's the means by which interaction or communication is achieved.

Most electronic systems have at least two interfaces: the thing producing functional results (e.g., a relay, printer, light, or display) and the method or organization defining functions to the operator or user.

The value of an interface depends more on the user's ability to comprehend its function, than on its actual capabilities.

WHAT DO WE EXPECT?

To be useful, the user must be able to perceive a device's function. The phrase "Form follows function" is particularly appropriate with a computer interface.

User interfaces generally fall into three types: intuitive, specific, and directed.

With an intuitive interface, the user needs no specific instructions to perceive its function or use it.

Generally, a push button is an intuitive interface. Think about elevators. It's usually obvious what is accomplished by pressing the button. Similarly, a doorknob or drawer pull is intuitive.

With a specific interface, function is easily understood with minimum information. If the push button is labeled "Exit" and is near a door, its application is easily understood by any user who understands the word "exit." Specific interfaces involve more of the user's time and concentration.

The directed interface gives specific instructions on how to operate the device. If a button reads, "To open door, press red button and hold for three seconds, then release," it would be considered a directed interface. A casual user who didn't read the instructions would probably push the button, try to open the door, and conclude the system was defective.

There are two principles to bear in mind. First, a device that requires no instructions is easier and faster to operate. Second, a device that requires instructions should be clearly labeled to avoid confusion.

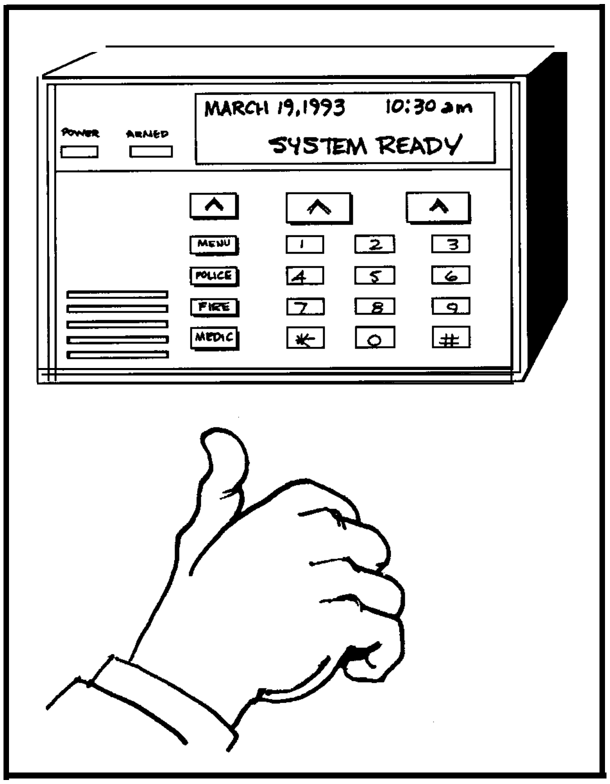

Figure *1: This alarm panel is accessible, but obtuse. What is the user supposed to do?*

USER FRIENDLY OR SEDUCTIVE

A user-friendly system is one that does not seriously challenge or obstruct the user-at least not as much as its predecessor did. The term implies that the user is able to use the device if the user chooses.

The user, however, frequently has a choice. If there is a simpler or more accessible means, the casual user chooses that one instead.

However, we want users to choose our products. A user-seductive interface is one that is not only usable, but desirable. It *leads* the user to perform the proper actions.

What makes a user interface seductive?

- Visibility-If a control is hard to find, see, or read, users are put off instantly. While this appears obvious, it is startling to see the number of products with LCD displays that cannot be read under normal (bad) lighting, can't be found at night, and have buttons that can't be read or are hidden behind doors.
- Clarity-Any user interface should clearly display what it is, what it does, and what its user options are. The user should not need a manual.

  Security controls are the worst offenders—do providers believe an inscrutable front panel is the first defense against intruders? If the user needs to enter a passcode, the device should read "Enter Passcode." On a complex display, buttons available for user input should be highlighted.
- Predictability-The device must respond the same way every time. Although multifunctional buttons conserve space and reduce cost, they are extremely confusing. Buttons should also respond the first time they are pressed.
- Feedback-The device should confirm that input has been received. Since accurate and immediate feedback is necessary, if confirmation takes longer than 100 ms, the device needs to provide a local indication to prevent the user from repeating the function, especially if the function is a toggle!
- Recoverability-The user device should be able to recover from errors gracefully, with specific prompts on how to proceed.

  Trapping users so they cannot return to the beginning, cannot undo an error, or must guess at options causes them to become hostile to the device.

  If the menu structure is multilevel, specific keys should return the user to the top level immediately or up one level per press. If the user jumps to the top level, it may be useful to provide a way to return to the previous place.
- Masterability-While menu-driven pages and extensive help are invaluable to a novice, they quickly become an obstacle for a power user. As a user gains proficiency, the system should provide shortcuts and macro functions to speed user input and provide a sense of mastery. Multiple paths let users find their individual comfort level.
- Customization-Function labels should correspond to their actual operation. The user interface should provide the ability to label buttons for each application. Ideally, the buttons should be easily relabeled aesthetically in the field.
- Serviceability-Service should be possible without inconvenience to the user or damage to the device or its mounting surface. Wall panels should be removable without damaging paint or wallpaper. If possible, routine maintenance should occur without removing the device or even remotely.
- Cosmetics-Most user input devices are installed in a variety of environments. The device should blend with most decor, have a variety of finishes, or be treatable to match existing finish. To cut cost and inventory, the visible portion of the device might be low-cost and separate. Then, custom finishing can be done before final installation.
- Comfort-The device should be usable for an extended period without causing the user discomfort. Size, weight, visibility, key pressure, feel, and location should be considered. The location for the device must be consistent with the nature of the interface.

Figure 2 illustrates many of these good characteristics.

INTERFACE TESTING

A good friend uses a "toaster test" to qualify a new technical product.

A toaster has the ideal interface. It has two slots about the same size and shape as bread and a single lever that moves in the direction you want the bread to go. There's no other user intervention, other than occasionally moving a small knob that's labeled "Light-Dark" or colored the different shades of toast.

While there is a manual-just to let you know not to wash the toaster while it's plugged in-it is superfluous. All you need to know is obvious.

All products should be this easy!

Many products in our lives are considerably more complex. Yet since we deal with them all the time, their operation is intuitive. A television is complex, but people know what they need and ignore other functions.

An automobile is an extremely difficult device. However, learning to drive is a common experience. Most people drive in a very short time. Once learned, it's as easy as a toaster.

An automatic teller machine is a more recent example of a learned interface. With the exception of a slot for cash, an ATM is not much different than a computer! The slot offers enough incentive that most people learn the operation quickly.

Moral: If your product dispenses cash, people will learn to use it.

While most folks easily operate their TV set, their VCR poses a more formidable challenge. The additional choices force the user to adapt to the machine's needs, rather than having the VCR respond to the user's desire.

Devices like VCR+ and on-screen TV guides greatly simplify the situation. Users can then say "Tape Married with Children" rather than "Record cable box channel 11 on next Wednesday at 9:00 P.M. for one hour at low speed." This sort of logic is alien even to some high-level COBOL programmers I know!

Also, the VCR gives no indication that it's been programmed correctly; you don't know until the show is over. This negative reinforcement causes many to be frustrated and give up.

A computer also is a masterable interface. With the proper software, it gives instant and positive reinforcement, so the user learns by practice. The software designer is responsible for keeping users out of trouble.

A security alarm system should be easily mastered. Its functions are limited, and there are few choices to make. Yet, many alarm controls are a source of frustration to their users. Many alarm manufacturers conceal the structure of the systems to confuse the miscreants. They only confuse the proper users, who daily use the system.

MASTERING THE INTERFACE

Because of our intelligence, humans take pleasure in mastering their universe. The pursuit of skill is a powerful motivator. The success of video games proves that people can be driven by no greater reward than the ability to become more adept at simple tasks.

The video game model can be applied to other interfaces as well. If your interface needs to be complex, challenging the user to learn its nuances is rewarding and results in greater comfort and familiarity.

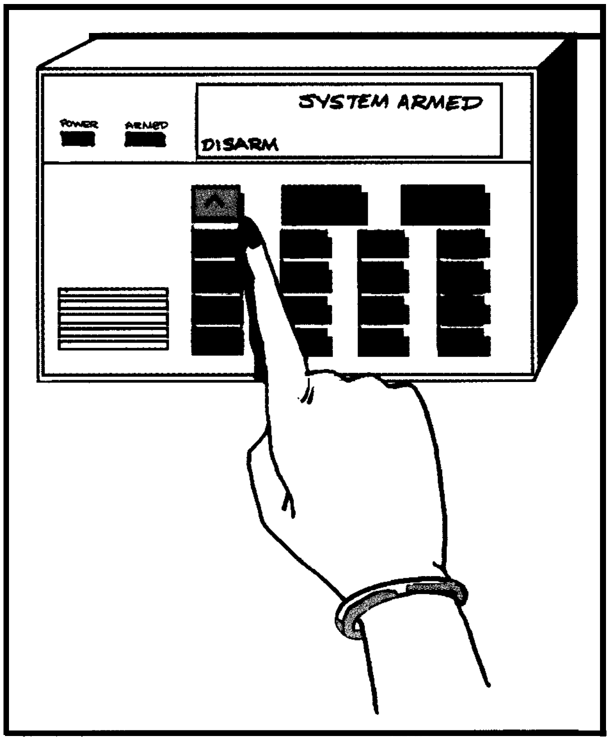

Figure 2: *This panel, similar co the first, has a crucial difference: the user is told exactly what to do. Presumably, after pressing "disarm," the user will be told to "enter passcode."*

Using rewards rather than error beeps to keep the user on-track is one method. Showing complex functions in a staggered manner after the user has mastered simple ones may be less confusing than presenting all functions at once. Group advanced operations separately from basic ones.

Functions may be accomplished by several methods. For novice users, a hierarchical menu structure with on-screen help avoids confusion. Shortcuts or macro keys aid the advanced user.

Let the expert customize the interface. Common functions should condense into a single keystroke. Rarely or never-used functions should hide or be removable.

It's useful to store both the customized and default function sets so either can be recalled. If the user can be identified by a passcode, the system could automatically load that user's preferences.

USER INTERFACE DEVICES

Switches are most suitable for obvious locations and basic functions. Access is fast and limited. Switches are generally mounted 3648" above the floor. Their function needs to be clear at a glance.

Keypads suit fixed locations and simple functions. Keypads without displays are generally 42–60″ above the floor, so functions should be clear and consistent. Keypads with displays should be placed at **eye** level (54–60″ above the floor).

Soft keypads, where the button function labels vary and are defined by the display, have more versatility and can be easily customized. As button functions vary with menus, visibility is critical. Because key locations are fixed, experienced users often press buttons faster than the display can respond. The system should buffer input to prevent the user from getting lost.

Touchscreens are extremely versatile and adapt to many input and display options. Because there's no tactile feedback, input tends to be slower than with fixed keys. Functions are frequently defined by menus, so locations may change as pages flip.

Touchscreens lend themselves more to menu-hierarchy functions than to data input. Users can offer rapid input more easily if expected sequences occupy the same physical coordinates.

Wireless remotes offer the convenience of freedom of movement. Most wireless remotes are one-way, so the user must rely on external feedback for proper operation. Power consumption must also be weighed against visibility, range, weight, cost, and operational life.

RF remotes offer nondirectionality, longer range, through-wall operation, and generally faster response. Backlighting should be strongly considered. For some systems, a docking station with battery charger may offer more reliable operation.

Using a TV set as a display device offers much of the flexibility of a touchscreen at a much lower cost. However, the TV has to be turned on to operate the system. While this limitation is not a problem for home theater or video-only systems, it may be inconvenient or intrusive for other systems.

Unless high-bandwidth displays can be guaranteed, display resolution also severely limits the usefulness of TV displays. Although interactive video lets a less expensive remote be used, parallax must be considered as the remote and display are not in the same field of vision. Devices like the "Air Mouse" may offer a solution if the eye-hand coordination issues can be resolved.

Telephones can be used as control devices. They are generally already in convenient locations, are familiar, and offer a variety of feedback options. However, the keyboard of a standard telephone set is too limited for a wide range of functions, and feedback generally requires the user to hold the receiver near the ear while waiting for a voice response.

Most people are only moderately familiar with the additional functions of an electronic key telephone set, and the displays of most sets cannot be accessed by control system designers. Furthermore, the location of telephones may not correspond to the desired location of input devices.

Voice recognition and response has been considered an ideal interface, as it requires little or no hardware or physical interaction by the user. However, it has many practical as well as theoretical limitations.

While faster processors, DSP devices, and cheaper memory have improved recognition systems, existing systems have severe limitations. User independence, continuous speech, context-sensitive processing, and inferred logic are difficult to achieve reasonably.

Psychologically, users aren't comfortable speaking to a system with low recognition rate and that appears to not have significant intelligence. Headsets make many people feel like they're talking to a wall.

AESTHETICS

For residential applications especially, aesthetics are fundamental to user acceptance. No matter how functional a device may be, if it offends the owner's taste or clashes with the decor, its applications may be compromised. It may be placed in an unobtrusive location, where its potential may go unexplored.

An integrated user interface replaces myriad single-purpose devices. This multifunctionality alleviates the "wall acne" caused by a wall surface being littered with a variety of aesthetically clashing and inconsistent devices. A common user interface generally appears more attractive and easier to use.

When integrating a number of systems, avoid ergonomic conflicts. Electric light switches are 36" above the floor, a convenient hand height. If the device contains a display, it's impossible to read at that height. Place the display at eye level—preferably the owner's eye level.

Placing the light switches at that height disturbs users who expect to find a switch just inside the door. You may need two interfaces or a keypad the user can find conveniently *before* entering the (dark?) room.

INTERIOR DECORATOR—FRIEND OR FOE?

Most individuals involved professionally with elegant homes have run into the interior decorator whose sense of style ended in the eighteenth century before electricity or even mechanics! They won't accept anything that isn't gilt, ornate, or draped in fabric, and they always want to place a huge painting wherever your keypads need to go.

Since most of their other clients apparently have servants, the notion of someone actually touching something in their rooms is abhorrent. They hide telephones in drawers, banish television sets, lock up remote controls, and generally make life difficult.

Usually, a compromise can be achieved. Get involved early and include their wishes in your design. They're frequently threatened by technology. Alleviating their fears may make them allies, not competitors. Just remember-your best referrals are from interior designers.

CONCLUSIONS

Much of this material may seem familiar and obvious-and it should be.

Surprisingly, though, most of the systems and devices we encounter in our daily lives intrinsically violate many of the principles discussed here. Measure and evaluate the effectiveness of your design against these human criteria.

Test your design with untrained users. If they can understand and effectively operate the device, it falls in the "obvious" category and will be easily accepted. If they are confused and frightened, redesign.

If they refuse to give it back, your design is truly user seductive!

Orrin Charm directs research and development for RRH Associates and is its user interface specialist. He is afield-trained expert on systems integration, user interfaces, wiring and installation practices, and documentation. He directed engineering at Interior Systems Design, helped found the Systems Integration Council of the Custom Electronic Design and Installation Association, and has served as its chairman. He may be reached at orrinc@ix.netcom.com or http://www.hometeam.com/orrinc/.

I R S

419 Very Useful
420 Moderately Useful
421 Not Useful

DEPARTMENTS

FIRMWARE FURNACE

Ed Nisley

Flat Surfaces, Widowmakers, and a Surfeit of Bugs

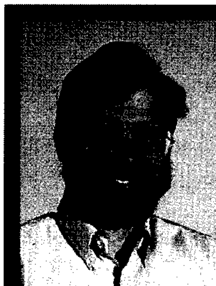

Ed covers two problems: an unsafe connector and a backup unit that won't back up. He walks us through his attempts to find solutions. His bottom line: don't cut design corners, think ahead, and be prepared for multifaceted problems.

Nature abhors a vacuum.
Benedict Spinoza

In my office, at least, nature also abhors a flat spot. 'Every surface normal to the local gravitational vector supports layers of interesting, useful, and essential widgets. When we began threading our way through heaps on the floor, something had to give.

So, I installed six bookcases with 120 square feet of new, flat surface area. They went in one at a time because Mary, who knows me all too well, gently refused to allow any clutter out of the office, even "for a few days while I get organized." I had to clear enough floor space to assemble each unit, clear enough wall space to stand it up, then repack it with good stuff.

In the process of touching everything in the office at least once, I came across a folder of column ideas and notes. I'll tuck them in here and there, whenever I have a bit of room.

In this issue, I'm including two interesting stories that may save your life and, perhaps, your hairline.

HOTWIRING A MODEM

Some years back, I briefly worked on a power-line modem project. Had things turned out slightly differently, it might have occupied the rest of my life.

When dealing with modems, it takes two to tango. The units on my desk, both development prototypes, bore serial numbers 3 and 4 carefully pasted on with Brady cable markers. As with most prototypes, they were pretty far out along the hand-filed aluminum axis.

Inside, the logic that would eventually fit into a 16-pin DIP sprawled out over two PC boards, with a huge pin-grid-array chip dominating the top board. Needless to say, custom PGA logic isn't cheap, but this was the fastest way to get the rest of the hardware and firmware checked out.

Power-line modems, by their very nature, have 120 VAC running among their innards. Each modem derived logic power from a 9-VDC wall-wart transformer, with a separate power-line cord for the data signal. An isolation barrier separated the high and low voltages.

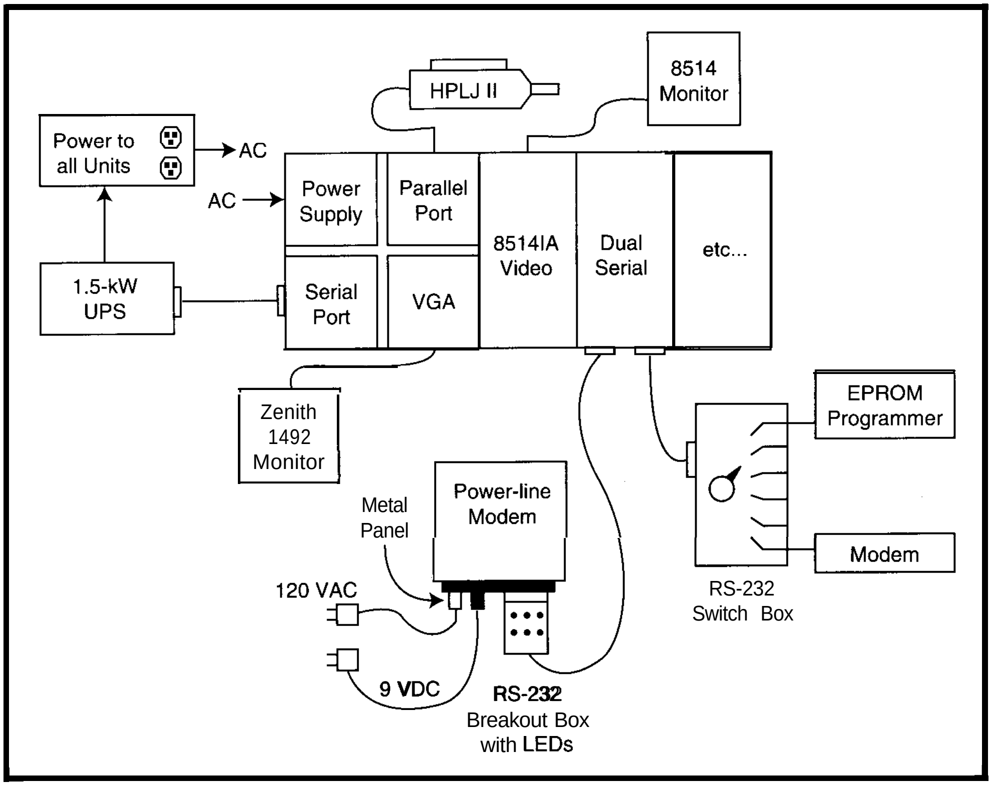

Figure 1—*Although the power-line modem circuitry isolated its logic circuitry from the 120-VAC line, the back pane/ a//owed a short-circuit connection between the power cord and the RS-232 Signal Ground. This figure shows some of the devices on my system*

Prototypes often omit safety features required in production units. For reasons that certainly made sense at the time, the line cord terminated at a female 3-pin DIN connector on the aluminum back panel. Thus, the line cord had a standard 2-blade polarized AC plug on one end and a 3-pin male DIN connector on the other. You'll sometimes hear engineers and technicians call a double-male line cord a ***widowmaker.***

The modems had demo firmware in their EPROMs to verify that they arrived undamaged. The transmitter sent a canned message, which the receiver echoed to its RS-232 serial output. I set the transmitter up in my machine shop, perched the receiver on the corner of my desk, plugged it into my PC, and fired up a comm program.

The received data had many errors, some of which appeared to be overrun errors inside the power-line modem. Thinking of the data and not the modem, I reached around to the back panel and pulled the AC line connector out between thumb and forefinger while keeping my eyes on the screen.

As the DIN plug came free, the "hot" pin brushed a socket-mounting screw, raising the entire back panel to 120 VAC. The mini phone jack for the 9-VDC supply has its ground side tied to the back panel, so the internal electronics followed the panel's potential.

Wall warts isolate low-voltage circuitry from the AC line. Ratings vary, but typical components break down around 600 VAC. Thus, raising the logic ground potential to 120 VAC shouldn't affect the circuitry's operation-unless, of course, a path exists from logic ground to power line neutral or ground, in which case you have a short circuit from the hot lead.

The modem case had plastic top and bottom halves that kept my wrist away from the back panel. The DIN connector's insulation held up well enough. For a split second, however, I held a fireball in my palm. However briefly, another path existed from the back panel to ground.

The modem's serial data passed through a metal-shell DB-25 connector. Although the RS-232 standard specifies pin 1 as Frame (or, sometimes, Safety) Ground, few devices connect that pin to the instrument's case. This modem used only the Signal Ground wire on pin 7, connected to logic ground on the circuit board.

The modem's RS-232 output went to a serial port on my IBM PS/2 Model 80- 111. That port connects pin 7 to system logic ground, which ultimately connects to the chassis, and from the chassis to the AC line's safety ground. If you have trouble following the connections, see the system layout in Figure 1.

I ever-so-carefully unplugged the widowmaker. The acrid smell of electrical death filled the air.

"DAMAGE REPORT, MR. NISLEY."

The Model 80, as you might expect, flatlined. Both monitors, an IBM 85 14 and a Zenith 1492, blacked out with their power LEDs still on. The HP LaserJet status panel read "Ready" and its On Line LED remained lit. The telephone modem was blank, as was my custom 6-way serial switch box.

Ever the optimist, I turned the Model 80 off and back on again. It produced

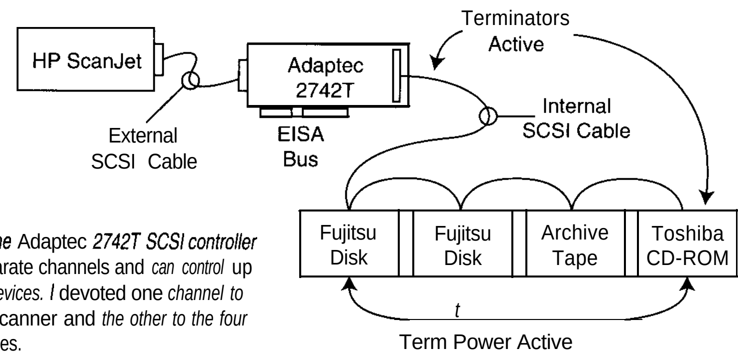

Figure **2**—*The* Adaptec *2742T SCSI controller* has two separate channels and *can control* up *to* 74 *SCSI devices. I* devoted one *channel to the external* scanner and *the other to the four* internal devices.

nary a power-LED blink nor a hard-drive clack.

I opened it up. The fan twitched when I flipped the switch, so the ("no user-serviceable parts inside") fuse inside the supply seemed OK. An IBM Dual Async card took the hit, so I pulled it out and tried again. After a pause, the power LED popped on. I closed it up, recabled the monitors, and flipped the switch. It continued to boot normally.

The Power-On Self Test noticed the missing Dual Async card and observed that the keyboard wasn't responding. I'd forgotten to plug it back in.. .

System diagnostics reported a dead system-board serial port, even though it passed the power-on tests. The suggested repair read "Replace System Board" because the Model 80 uses surface-mount ICs.

That port connected to the 1.5-kW UPS driving the entire system. Obviously, the UPS remained functional, but would I ever talk to its microcontroller again?

Everything else inside the 80, including **12** MB of RAM, an 80387, and a 1 15-MB ESDI disk, seemed OK. OS/2 booted normally, so the hardware wasn't too badly dinged.

Although I'm sure the Dual Async card was repairable, I opted to replace it. A local store had one in stock for about $200. Hey, Microchannel seemed like a good idea at the time.

The repair tech said that a PS/2 Model 80-l 11 system-board exchange ran just shy of $800, but they might waive the labor charge because the part (singular!) is so expensive.

My six-way RS-232 switcher was plugged into the other port on the Dual Async card. It routed the catastrophe into my EPROM programmer, blowing away its 1488 and 1489 serial-interface chips.

Both the phone modem and UPS survived unscathed. I'd left the modem turned off, which accounted for the blank display. The UPS evidently didn't take a direct hit. The LaserJet, on a parallel port, kept on ticking.

The MAX232 chip inside the power-line modem popped. The rest of the modem's circuitry, including that precious PGA, seemed unharmed.

The smell of electrical death came from an RS-232 breakout box plugged directly into the power-line modem's serial port at the back panel. A PC board trace connected the serial logic ground pins; the fireball vapor-deposited copper plating over the case interior. I soldered a length of wire-wrap wire between the two pins, leaving the smudge and smell as reminders.

I think the trace acted as a fuse, saving the rest of my system from even more damage.

THE MISSING LINK

A fundamental principle of accident prevention goes like this-it takes two people to make an accident.

The first person sets up the conditions, either deliberately or inadvertently. The second person doesn't notice the situation and triggers the accident. In some cases, one person can set up an accident and then walk right into it. While that may not always be true, if you recall the accidents you've survived, it covers many of them.

Ken and I noticed those cords when we unpacked the modems at Micromint world HQ. At the time, we shook our heads over an obviously dumb idea. I thought of replacing the connectors or gluing the cords into the DIN sockets, but never followed through.

The modem's internal wiring exposed a lot of metal at 120 VAC. The folks who put it together thought that their users would see only a closed box, with no need to fiddle with connectors, change EPROMs, or do anything dangerous. They're going to do a much better job on the next prototypes, you can be sure.

Yes, I accept my share of the blame lying around this accident. Remember my story while building your next gadget. If you wonder whether you should take that extra step to make it a little safer or a little more convenient, just do it! Remember it when you work around lethal voltages.

Next time, some body might provide the ground path.

A DRIVE CRASH

Shortly after moving to Raleigh, I replaced that venerable Model 80 with an EISA 486DX2 system and discussed the details in my *INK 47* column. The system ran fine for months, then developed a curious problem.

My system has a dual-channel Adaptec 2742T SCSI-2 adapter for the hard drives, CD-ROM, scanner, and an Archive Viper 2525S backup tape drive using 525-MB, ¼″ cartridges. Figure 2 sketches the SCSI chains.

Everything worked fine, except that some tapes crashed the drive during its loading sequence. The drive normally buzzed, clicked, and whirred while aligning its heads to the tape tracks. Loading a bad tape produced a click, a buzz, a whir or two, and then the drive halted. The sole LED on the drive's front panel began blinking: two blinks, pause, four blinks, pause, and one blink.

You've seen techniques for generating catastrophic error codes like that, right here in this very column.

The drive remained dead to the world until I reset the whole system. Attempting to access the drive hung the CPU because the SCSI controller jammed waiting for a response from the drive. The red Drive Activity LED stayed on forever.

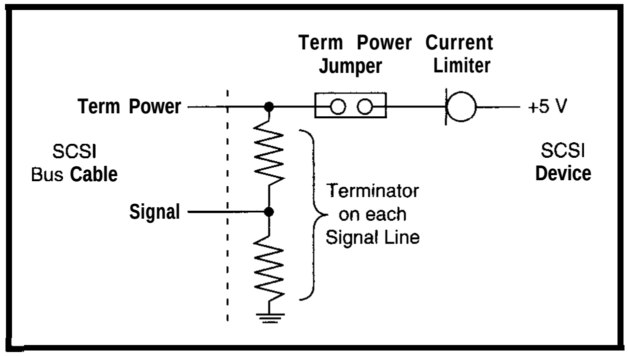

Figure 3—**Each** *SCSI* **device** *can provide power for the bus* **terminators through** *the TermPwr pin. Although* **only** *one device should power the bus, I* **inadvertently** *left three TermPwr* **jumpers installed.**

Want to sound like Mission Control? Repeat after me: "Mmmm, not a nominal outcome."

Although a bad tape shouldn't cause a complete drive lockup, I figured maybe I got a bad batch of 3M tapes. They have a lifetime warranty, so I returned one to 3M for analysis. According to 3M, the tape was well within specs. They bulk-erased it, reformatted it, retested it, and it worked fine for them. When I popped it in the slot, the drive crashed again.

SCSI devices have a notorious sensitivity to bad cables, improper termination, and suchlike. I moved the drive to a different connector, built a new cable for the 2742T's other channel, then tried a variety of tweaks. Nothing helped. Bad tapes remained bad and the drive continued to work fine on the remainder.

Thinking I might have a defective drive, I sent it back with a bad tape for a checkout. As you might expect, both passed all tests with flying colors. Once here again, they failed again.

By now, you're thinking, "It's the controller, stupid!" That's what I thought, but the folks in the shop checked the drive with their 2742T. Same drive, same model of controller, same tape.

CHIPPING AWAY

A year later, half a dozen cartridges sat on the shelf, each marked "Bad Tape!" in red ink. At 30 bucks a pop, that represented a major motivation to solve the problem. When the drive crashed on another tape, I decided, at long last, to get serious.

I sent three tapes back to 3M, including the one they'd seen before. They reported that all were in good shape, but replaced the old one on general principles. That tape worked in my drive, while the other two still failed.

Conner Peripherals had since bought out Archive. Their tech support folks decided a firmware crash like that certainly indicated a drive problem. Unfortunately, swapping my drive for a refurbished unit cost 300 bucks, as it had but a one-year warranty. I sighed, wrote the check, and sent the drive back again.

The refurbished drive failed exactly the same way.

I called back and the tech asked about the EPROM version. He decided that, perhaps, upgrading from Version 006 to 007 might help, but he couldn't send one out directly. After an afternoon on the phone with Conner's Tech Support, Customer Service, Express Parts, and other voices too numerous to mention, I finally found someone who could send me the EPROM.

With Version 007 installed, the drive successfully loaded all the previously "bad" tapes. I fired off a backup that wrote 480 MB of data to the tape. Writing seemed to work fine.

The backup software rewound the tape, started its AutoVerify pass, and halted with an "Unrecoverable Tape Error" after reading about 250 MB. On other tapes, it reported "Cannot Write Tape" after successfully writing the fileset header. Upgrading the backup software from Version 1.11 to 1.13 didn't change the symptoms.

I repeated the SCSI termination, cabling, rechannelling, and twiddling to no avail. Finally, I called Adaptec to discuss the situation. Their tech suggested a few setup options: limit the drive's SCSI data rate to 5 MBps, disable synchronous negotiation, and reduce the controller's data FIFO usage to 75%.

He also sent me a current BIOS EPROM, gave me a Case ID that would bypass the telephone queue if I needed further assistance, and reminded me how to use their phone system.

None of the adjustments had any effect. But, with the new Adaptec EPROM installed, a backup using the same tape ran to completion. The AutoVerify step reported that several files had verify errors indicating that the data on tape didn't match the original file.

I restored the files to another drive and compared them with the originals. Each file had one or two mismatching bytes, in places where the original file was obviously correct. Restoring the files several times produced the same results. Either the files were written incorrectly or the restores failed the same way every time.

When I attempted to call Adaptec, their phone system announced that all the menu options had changed. Nowhere in the new structure could I find the "Enter your Case ID number" option. I hear that Future Domain bought Adaptec and the phone-system changes reflect a new approach to tech support.

After several attempts over a few days, I finally gritted my teeth while spending half an hour on hold. Amidst peppy hold music and ads for their $35per-call priority support, the phone system asks for your phone number. When the tech got around to me, he had my complete record at the ready without reference to the Case ID number.

We reviewed the situation and shook our heads.

I mentioned that my 2742T board sported hand-soldered components atop several chips and Engineering Change wires tacked down with hot-melt glue. He flinched audibly, agreed a board swap might be in order ("just to get that board out of circulation"), suggested several additional tweaks, and offered to call me back in half an hour.

SCSI terminating resistors can draw power from either the bus cable or the device. A jumper in each device connects logic power to the TermPwr bus line and from there to the terminators. Figure 3 shows some typical circuitry. In my system, the terminators received power from the two disk drives and the CD-ROM drive in addition to the 2742T controller.

Obviously, connecting two or more voltage sources to the same wire smells like a Bad Idea. Reviewing the manuals for the drives and controller showed I hadn't overlooked a warning. To quote one description: "JP12 Alternate Terminating Powers: When a jumper is installed, +5-V alternate terminating power is available at pin 26 on the SCSI interface connector."

Sounds harmless, right?

The tech said this situation often causes problems in network servers, multitasking systems, and other applications with lots of bus activity. If any of the cross-connected power supplies glitch during a bus transfer, the whole system stumbles. Even if a device doesn't have its terminating resistors installed, as was the case with my disks, the TermPwr line can still cause trouble.

Removing the extraneous jumpers didn't eliminate the errors, however. Installing the new 2742AT controller did, but the tape continued to fail about halfway through either the backup or the verify step with mysterious errors. It failed the same way on the secondary channel, with a different cable, and, yes, with a variety of tapes.

Having replaced everything else, I returned the tape drive once again. Much to everyone's astonishment, this drive actually failed in the shop by reporting a blank tape halfway through the verify step. It had other troubles as well, which led all of us to wonder how it got past Conner's final QC inspection.

When I started this column, I planned to report a successful conclusion after replacing the tape drive for the second time. Regrettably, the box I got yesterday contained a drive with bad motor bearings (I wore earplugs while making the test backups) and firmware that crashes with blink codes 1-1-3 and 2-2-7.

Back to Square Zero.

HERE A BUG, THERE A BUG

Most problems have one root cause that, once you figure it out, resolves all the issues in one fell swoop. Sometimes, though, the symptoms have nothing to do with the problem. Other times, you'll find many problems lurking under each rock.

In this case, I believe my system had several incompatibilities that I worked through in sequence. The original Version 006 firmware had a definite problem, as evidenced by the fact that it crashed with no intervention from any other code. Something similar applies to Version 007.

The tapes contributed to the confusion, as some tapes failed while others worked correctly. Were I using surplus tapes from Bottom Feeder Magnetics, such behavior might be understandable. Conner/Archive had 3M tapes on their Approved list and the folks who sold me the drive use them regularly. Perhaps something happened to all two dozen of my backup tapes!

Until I installed the new drive, I never had trouble with random byte errors. However, the Adaptec 2742T's hand-wired ECs were applied to the Dallas DS2107S SCSI Terminator chips. I suspect the new 2742AT's circuitry has more oomph or stability. In any event, the new tape drive exposed a problem I didn't think I had.

Next step: a 2.5GB drive.

RELEASE NOTES

I can't think of anything to post on the BBS this month! Next month, we'll take a microscopic look at 80x86 performance. ❑

Ed Nisley (KE4ZNU), as Nislep Micro Engineering, makes small computers do amazing things. He's also a member of Circuit Cellar INK's engineering staff. You may reach him at ed.nisley@circellar.com or 74065.1363@compuserve.com.

IRS

422 Very Useful
423 Moderately Useful
424 Not Useful

Stepping into an Analog World

FROM THE BENCH

Jeff Bachiochi

Most analog control doesn't have to be elegant, although it can be with PWM. Here, Jeff adds a twist to combining analog and digital in a motor-control system—he controls digital stepper motors with an analog input voltage.

Today most battery-operated vehicular platforms use simple DC motors. Whether they speed around an RC raceway or wind their way through a robotic maze, analog control doesn't need to be elegant. However, by adding PWM control of the motor (a touch of digital if you will), it can be.

This month, I want to add a twist to the art of combining analog and digital into a motor-control system. I'll show how to control digital stepper motors with an analog-input voltage. With my control voltage of O-5 V, 0 V is the maximum reverse speed and 5 V is the maximum forward speed.

Why choose a stepper motor?

In my case, I have a slew left over from my column on disk drives (INK 51). I admit PM DC motors (especially in the surplus catalogs) are easier to use, but many run at very high RPM and are not applicable to robotics unless there's a gear-reduction system.

If you stay within the torque curve of the stepper motor, you can calculate distance by the number of steps you take. Once outside this curve, missing a step degrades your ability to estimate distance accurately.

STEERING AND MOTION

I'd like to keep this platform as small and light as possible. To do this, I use two motors, one for each drive wheel. These motors are located at the three and nine o'clock positions on the base (see Photo 1). An idler wheel takes up the six o'clock position, keeping the platform level.

The four-step pattern of most stepper motors is illustrated in Figure la. As each coil is energized, the rotor is pulled one step clockwise or counterclockwise, depending on the sequence.

A half step is created when two coils adjacent to one another are energized at the same time. It can be accomplished by using the sequence shown in Figure I b. Both coils pull on the rotor forcing it to remain in between the quadrature step positions.

By using these half-step positions (i.e., the points where two coils are energized at once) as quadrature stations, I get an increase in overall torque (see Figure lc). The increase in torque remains as long as two coils are always energized at the same time.

For the most part, the steppers found on old disk drives (and obsolete pin printers) are unipolar. The four coils have a common lead or sometimes two, one for each opposing pair

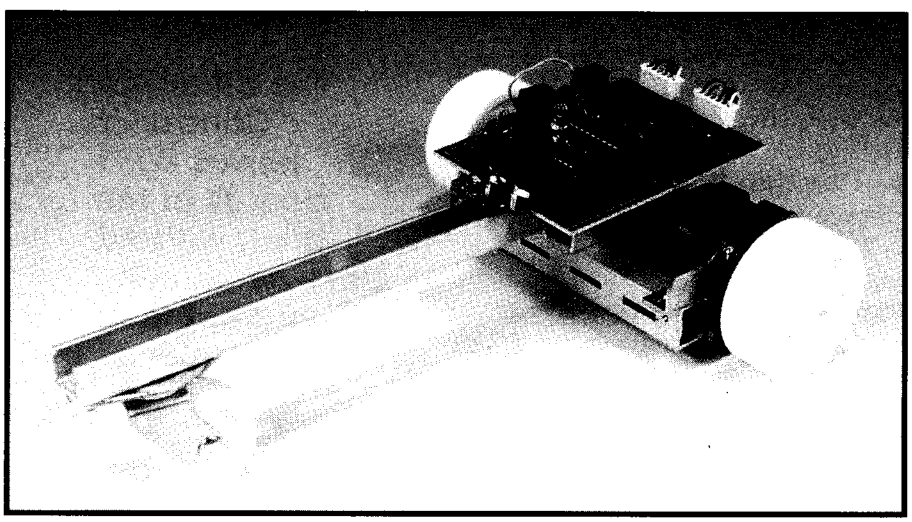

Photo 1—*The mobile platform is made from lightweight aluminum wall-bracket girders, with chair casters for wheels.*

of coils. You only need to apply a single-polarity voltage to the coils in a progressive fashion to get the rotor to turn.

A bipolar stepper motor has four leads and is similar to what is necessary to reverse a DC PM motor. Drive polarity must be swapped on each set of coils to achieve rotation. Although these can be used for this application, I won't discuss them.

To drive the four coils of the unipolar stepper motors, I use open-collector current sinks. Eight drivers are needed for the two motors. Individual transistors could be used, however the ULN2803A octal-drive package has all the necessary transistors in a single 18-pin DIP.

I can get up to 500 mA of drive current from this package using up to 35 V (just keep an **eye** on the total package dissipation). My steppers draw only 200 mA each.

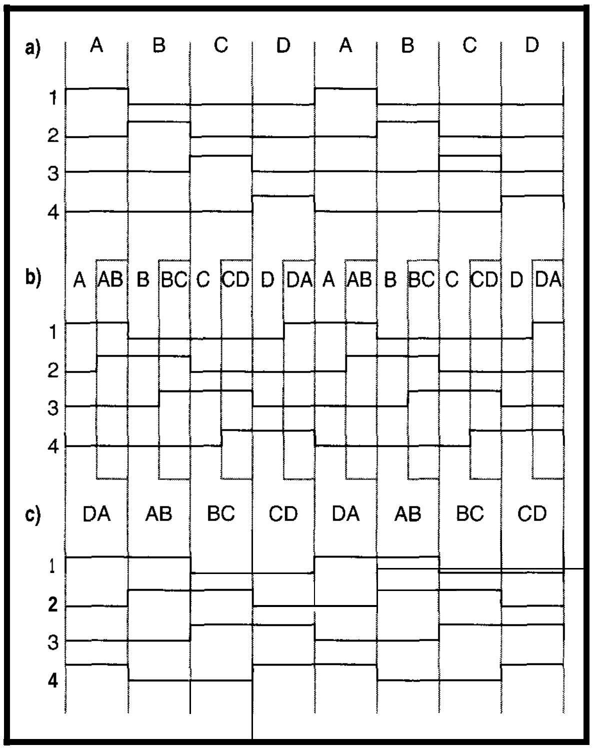

Figure 1—Stepper motors can be made *to behave* **differently depending** OR *how they* **are driven. Three** *stepper motor* **phase timings include (a) single phase, single step,** *(b) half* **step, and** *(c)* **double phase, single** *step.*

In addition, the '2803 has internal protection diodes across each coil to reduce EMF spikes. Control inputs for each transistor are TTL compatible, which eliminates the need for series resistors, and simplifies the interface.

SEQUENCING THE STEPPERS

Stepper-motor controllers are available to provide the sequencing necessary for rotor rotation. These controllers require a direction and step input. By using a small micro to take the place of a stepper controller, we can simplify the front end of this project by letting the micro control an A/D converter.

Some of the newer micros are available with built-in A/D converters and come close to the price of a PIC 16C54 and an external ADC0832 (see Figure 2).

Most internal A/D converters are limited to 8 bits of resolution, so using an external A/D converter gives you the option of 8-, 10-, or 12-bit resolution. The least-expensive PIC, the 16C54, can easily read both of the channels of an ADC while mimicking the outputs of two stepper-motor controllers.

Each motor requires four output signals. The quadrature states for these outputs are labeled A, B, C, and D. Four bit-set or -clear commands are used for each state. Look again at Figure 1c. Notice that only two output changes are needed to shift to a new state. One coil needs to be deenergized while the opposite coil needs to be energized.

I chose to perform a function on each of the four outputs whether they actually needed to change state or not for two reasons. First, this simplifies the coding so the same routine can be used independent of the direction of rotation. The STATE routine always sets all four outputs properly.

Second, if by chance the program flow makes a wild jump to an out-of-sequence state, not more than two coils can ever be energized at the same time.

SPEED AND DIRECTION

The analog input voltage is converted into an 8-bit value from 0 to 255. The micro checks the most-significant bit of this value and determines which direction the motor should turn. Values over 128 are considered forward movement and those less than 128 reverse movement.

A direction bit for each motor indicates whether the state should be incremented or decremented. The new state of the outputs is changed based on the remaining 7 bits of the A/D conversion.

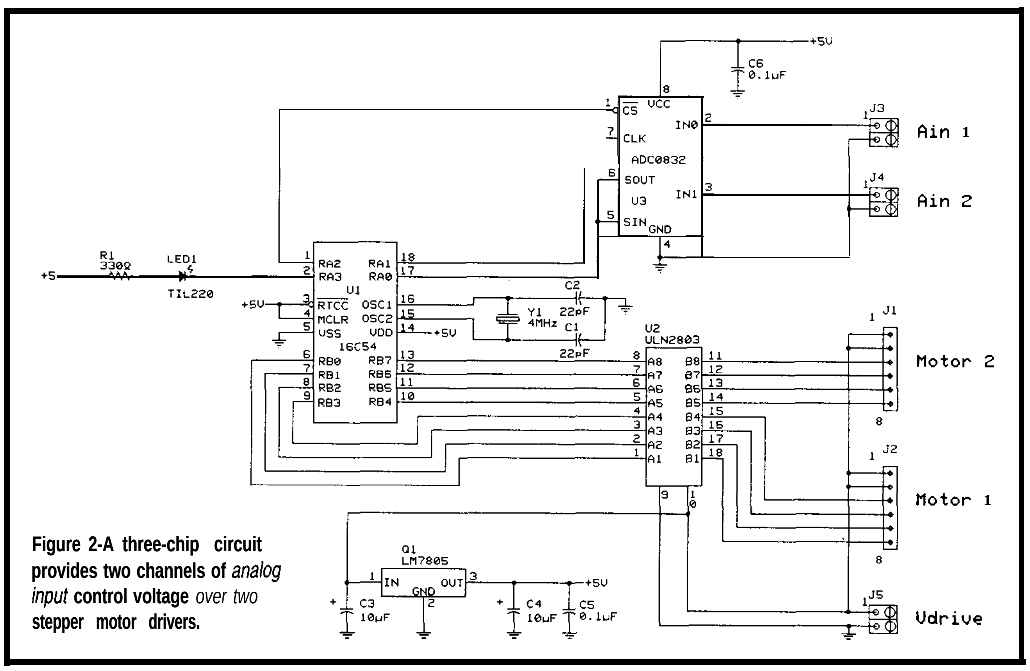

Figure 2-A three-chip circuit provides two channels of *analog input* **control voltage** *over two* **stepper motor drivers.**

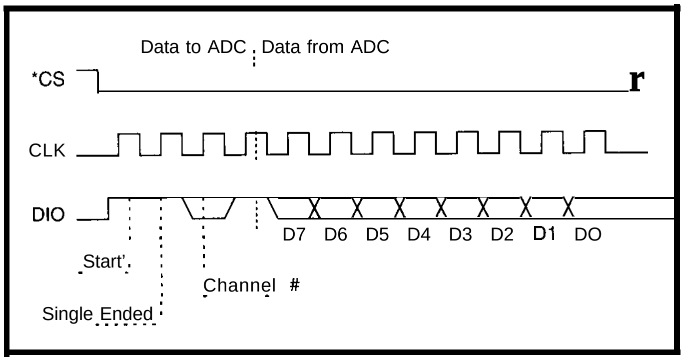

Figure 3—This serial sequence provides the A/D conversion.

After determining the direction of each motor, the micro complements the conversion values when the direction is forward (MSB set). If you refer to Table 1, you'll see what happens when the conversion values are complemented based on the MSB.

As the input voltage moves either above or below 2.5 V, the computed values become smaller and smaller. It is these values which are loaded into countdown counters. Input voltages close to 2.5 V produce long countdowns and inputs to either extreme produce short countdowns.

The time between steps (or states) is this countdown period. To adjust the maximum step time (minimum countdown), a delay of one TMRO rollover is used. By choosing a prescaler for TMRO, you can choose to add binary increments of 1 ms to each count.

The minimum step rate (i.e., maximum count) is never infinite. So, there is always movement, even if it is relatively slow. To correct for this, the micro tests the computed value for the maximum countdown time (0111-1111). If it is found, the countdown and step routine is skipped.

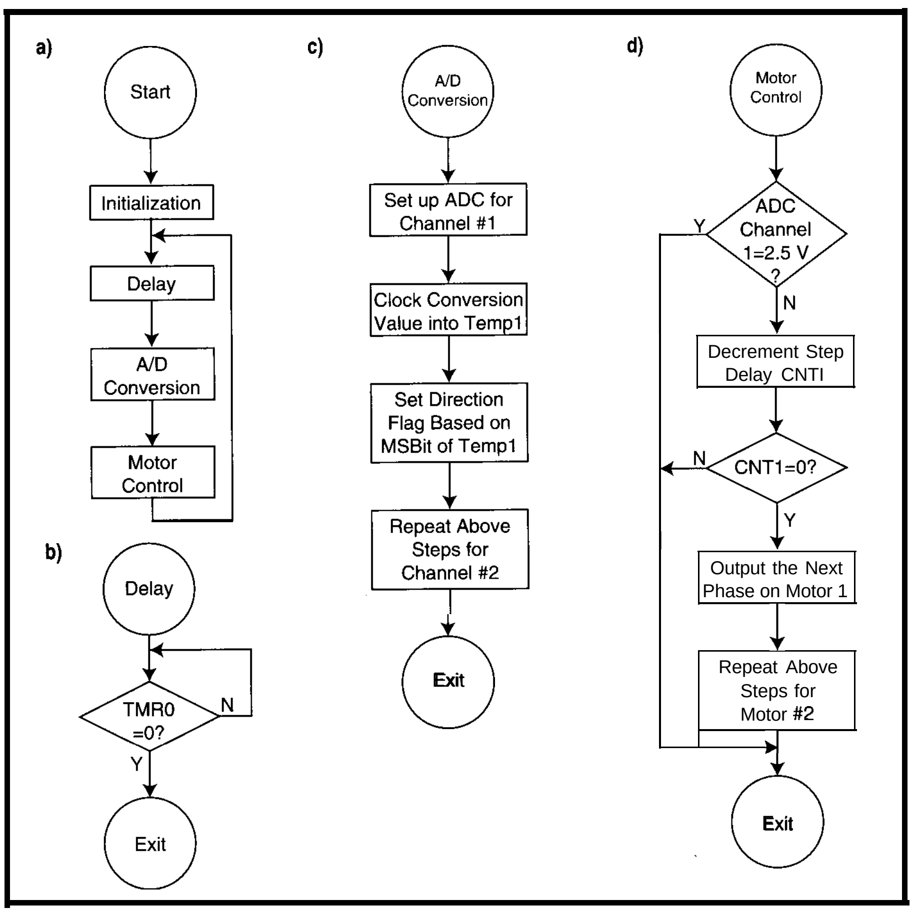

Figure 4—After initialization (a), three main routines are used in the application: (b) delay, (c) A/D conversion, and (d) motor control.

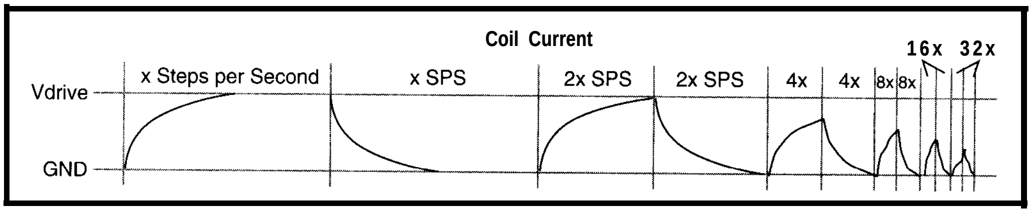

Figure 5—*Notice how the coil current is affected as the step time is reduced to less than the current's rise time.*

A/D CONVERSION

Serial A/D converters are very popular these days, and for good reason: they are small and inexpensive. The ADC0832 easily falls into this category.

Since communication between the micro and A/D converter is serial in nature and the '54 doesn't have an SPI port, a touch of bit twiddling is necessary. Three I/O bits (● CS, CLK, and data) are used.

If you do not have bidirectional I/O bits, the A/D converter's Din and Dout lines can be used independently. Otherwise, as in this case, they can be tied together.

The '32 requires three bits of setup data: a start bit (1), the channel configuration (single ended/ differential, O/l), and the channel to convert (O/l). These are clocked in by the micro's CLK output. Following an extra clock cycle, the A/D converter places the converted data on Dout a bit at a time in sync with your clock.

Data bits to the A/D converter must be latched before the rising edge of the clock and returning data bits read after the rising edge of the clock. The sequence for reading the A/D converter is shown in Figure 3.

CHARTING THE FLOW

You can follow the program flow by referring to Figure 4. Only three routines are necessary in the main program loop for this simple interface.

Power-on reset begins with the normal initialization sequence followed by the main program loop. As Figure 4a illustrates, this loop consists of a delay, an A/D conversion routine, and the motor-control code.

The delay-a rollover of TMRO-is set by choosing a prescale divider which sets the maximum step rate of each motor (see Figure 4b).

In Figure 4c, the A/D conversion routine reads one channel of the A/D converter and assigns the conversion value to variable **TEMP1** and direction flag **FLAGS. 0** based on the most-significant bit of **TEMP1.** The second half of the routine does the same for the second ADC channel and assigns the conversion value to variable **TEMP2** and direction flag to **FLAGS. 1.**

In Figure 4d, the stepper motor routine decrements CNT1 each time through and exits if it's not equal to zero. CNT1 is initially loaded with **TEMP1** each time it reaches zero. The smaller the number loaded into CNT1, the fewer times it loops before reaching zero.

Approx. Input Voltage	Conversion Value	Direction Flag	Computed Value
0.00 v	00000000	R	00000000
0.10 v	00000001	R	00000001
2.48 V	01111110	R	01111110
2.49 v	01111111	R	01111111
2.50 v	10000000	F	01111111
2.51 V	10000001	F	01111110
5.99 v	11111110	F	00000001
5.00 v	11111111	F	00000000

Table 1--The *code converts the* **analog** *input voltage to direction* **and** *speed* **data.**

When it does reach zero, the rest of the routine "kicks" into phase. The direction flag determines whether the phase is increased (forward) or decreased (reverse).

The new phase directs the flow to one of the four possible motor phase output tables. Each table sets or clears each of the four motor outputs associated with each individual phase (A, B, C, or D).

The second motor is handled the same way. Finally, execution jumps back to the delay loop for another cycle.

SPEED LIMIT 55-NOT

As Figure 5 demonstrates, maximum step speed (and torque) become a concern when the step rate is reduced to a point where the coil current cannot rise sufficiently before falling due to the motor's inductance.

Using a higher drive voltage and constant-current supply (instead of the ULN2903A) could provide a higher maximum stepping rate. However, higher voltages are generally not available on robotic platforms, so I'll be satisfied here with a fixed maximum of less than 500 steps per second.

By far, the most powerful stepper I've come across is the carriage drive from a dot matrix printer (0.9 A at 4 V). When choosing to use a higher-power motor, make sure the drive circuitry can handle the higher currents.

A tweak on this system might be to replace the A/D converter front end with PCM input-you know, the signal stream used to control RC servos.

Are you asking why I didn't choose that kind of interface right from the start? Well, it's a secret for now, but here's a little hint, "Not all computers interface digitally." ❑

Jeff Bachiochi (pronounced "BAH-key-AH-key") is an electrical engineer on Circuit Cellar INK's engineering staff, His background includes product design and manufacturing. He may be reached at jeff.bachiochi@circellar.com.

UPDATE

The power-line modem chip used in *INK 64-65* can be ordered now from Mouser Electronics (part number ST7537HSICFN, $17.64 in single quantities). Call (800) 3466873.

I R S

425 Very Useful
426 Moderately Useful
427 Not Useful

Pick a Peck of PLDs

SILICON UPDATE

Tom Cantrell

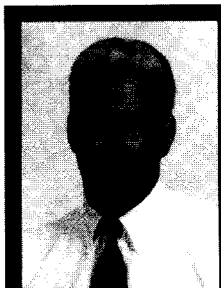

Believing PLDs to rate a spot in the micro-electronics revolution right up there with heros such as the microprocessor, DRAM, and EPROM, Tom reviews programmable logic ICs. His focus: Cypress Semiconductor.

When it comes to the microelectronics revolution, no one can deny PLDs deserve a spot in the hall of fame right next to heroes like the microprocessor, DRAM, and EPROM.

OK, pop-quiz time. Who comes to mind when you think of major PLD suppliers? I'd say the list includes AMD, as holder of the original MMI (Monolithic Memories) crown jewels (PAL patents); Xilinx as the FPGA pioneer; and a list of PLD specialty shops like Altera, Lattice, and so on.

Your list, like mine, probably does not include Cypress Semiconductor. Frankly, my awareness of Cypress was more related to the outspoken antics of founder T.J. Rodgers than their product line.

Simply put, when it comes to politics and management philosophy, T. J. is no Mr. Rodgers. His periodic broadsides in the editorial pages of the local newspaper are something to behold.

A good example was his recent "Litigation Reform: Did Congress Do the Right Thing?" *(Sun Jose Mercury* 12/24/95). He relates the story of how Cypress, like many other Silicon Valley companies, was sued for "fraud" when their stock dropped.

T.J. first politely points out that one of the purportedly injured investors had been similarly "defrauded" **18** times before and might consider a new line of work. He then makes a cautious and reasoned response to the legal threats: "The plaintiffs will get their first nickel out of me when they pry it out of my cold, dead fingers."

Anyway, I-as you may be-was surprised to discover that Cypress actually offers quite a selection of programmable logic ICs. So, let's get off the op-ed page and get into the data book to see if T.J.'s chips are as hot as his opinions.

FAST PALS AT SILICON VALLEY HIGH

By now, everyone should be familiar with the good old PAL architecture, whose basic sum-of-products architecture is little changed since its invention by MMI in the '70s.

As shown in Figure la, the inputs feed AND gates (i.e., the product) whose outputs feed an OR gate (i.e., the sum). The PAL is programmed by deciding which inputs feed an AND gate by making or breaking the appropriate connections.

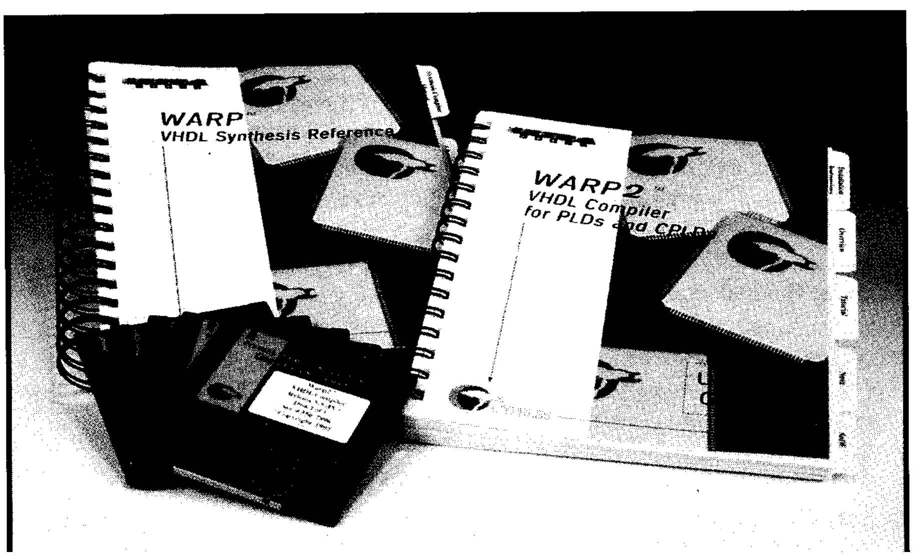

Photo 1—*At only $99,* **WARP2 eliminates** *the sticker* **shock** *traditionally* **associated** *with* **high-end** *EDA fools like VHDL.*

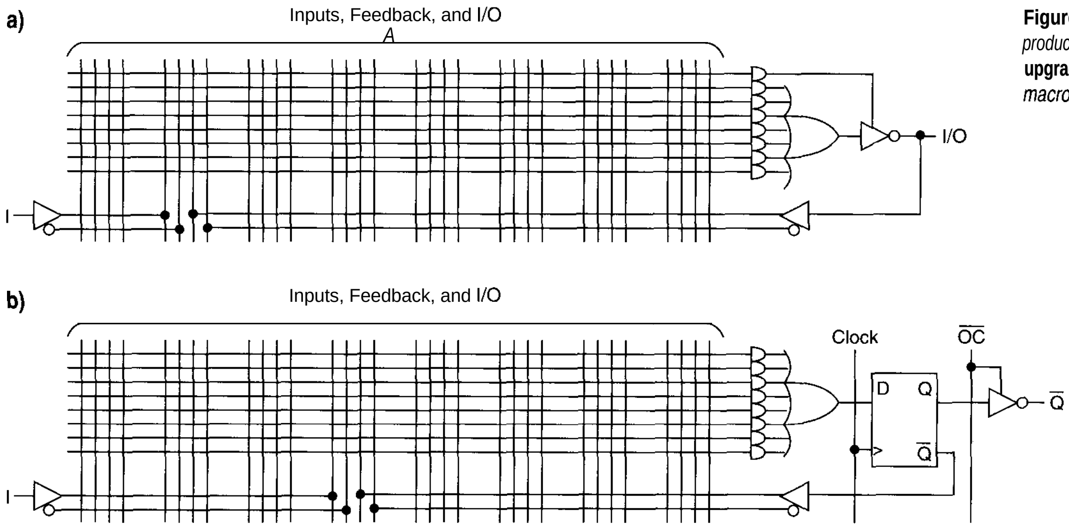

Figure I--The *original PAL sum-of-products architecture (a) has* **been upgraded with evermore** *complex macrocells (b).*

Though the PAL's basic sum-of-products architecture hasn't changed, architectural innovation has occurred on the I/O side with the simple output gate replaced with evermore complex macrocells as shown in Figure lb.

Consider the macrocell of the popular 22V10 (see Figure 2a) shown in Figure 2b. As you can see, it offers a selection of combinatorial or registered logic with high or low polarity. As well, the pins can be designated as input or output.

A key byproduct of smarter I/O is that a single device can be programmed in a number of permutations that originally each required a dedicated part. The chip's 22 I/O lines can be configured anywhere between having up to 21 inputs with a single output or as 12 inputs and **10** outputs.

Another nicety is the addition of common preset (synchronous) and reset (asynchronous) lines that eliminate the need to dedicate product terms for those functions.

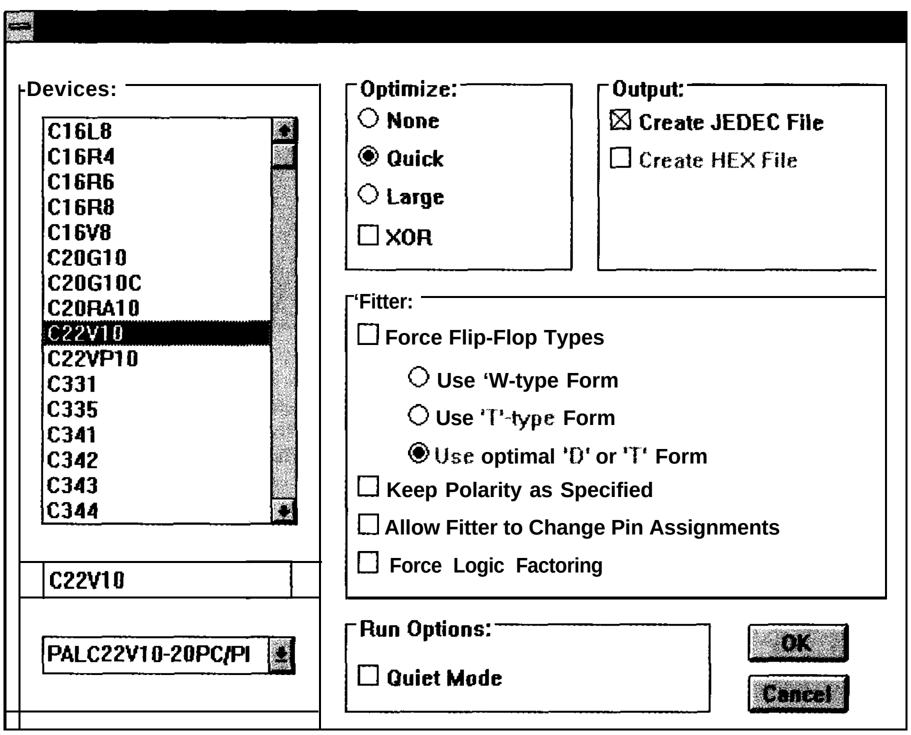

Photo 2—*Thanks to the fixed layout and* ***predictable*** *timing of PLDs, designers have only a few WARP2 compile options to deal with.*

Like others, Cypress is quickly moving away from bipolar fused technology to CMOS flash-based reprogrammable devices. However, what differentiates Cypress PALs from most of the others is speed.

For instance, the 22V10 D-7 version achieves rather astounding specs. It sports 7.5-ns propogation delay and maximum frequency (dictated by clock setup, hold, and output delay specs) of 100 MHz (state-machine, external feedback), 133 MHz (state-machine, internal feedback), and up to 166 MHz (no feedback, i.e., datapath only).

Yes, there's a price to pay-$5.95 in volume-and it sucks power like Las Vegas (140 mA @5 V), but otherwise it's just the ticket for homebrewing your own Cray!

SEA OF LOGIC

There's nothing really that complex about so-called CPLDs (Complex PLDs). Just take a bunch of PAL-like logic blocks and stick them on a single chip with some kind of interconnect scheme to hook them together.

A good example is the Cypress Flash370 family, depicted in Figure 3, that combines a number of logic blocks (2, 4, or 8, depending on the version) with a Programmable Interconnect Matrix (PIM). Volume price is about $2 per logic block; not bad since each of them pretty much matches the features, speed, and capacity of any PAL.

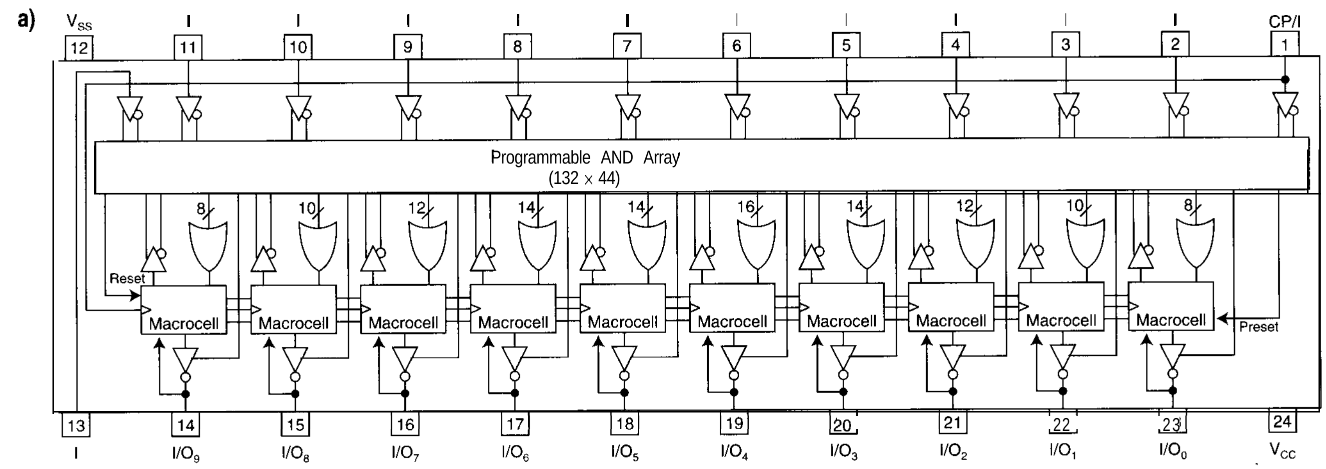

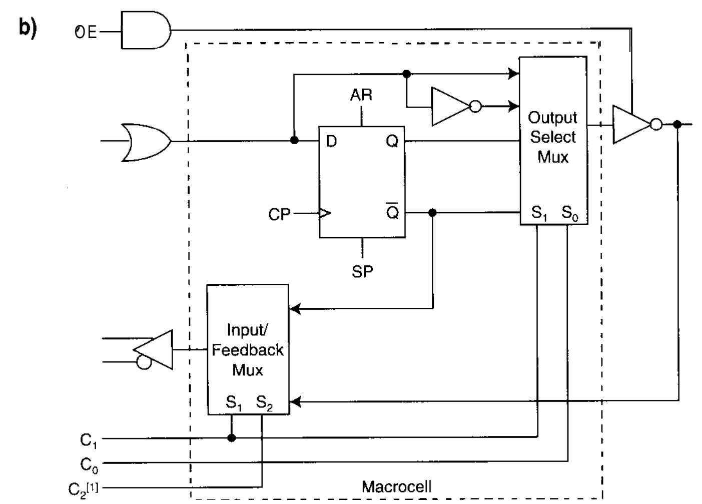

Key:
AR = Asynchronous Reset
SP = Synchronous Reset
OE = Output Enable
CP = Clock Pulse

Output Macrocell Configuration

C_2[1]	C_1	C_0	Output Type	Polarity	Feedback
0	0	0	Registered	Active Low	Registered
0	0	1	Registered	Active High	Registered
X	1	0	Combinatorial	Active Low	I/O
X	1	1	Combinatorial	Active High	I/O
1	0	0	Registered	Active Low	I/O[1]
1	0	1	Registered	Active High	I/O[1]

Note: PAL22VP1 OG only

Figure 2-One **of the most popular** *PLDs is the 22V10 (a), largely thanks to the versatility of the I/O macrocell (b).*

As shown in Figure 4, the product-term array accepts 36 inputs from the PIM (spawned from feedbacks or pin inputs) along with their inversions, and it generates 80 general- and 6 special-purpose (4 output enables and 2 resets) outputs.

The outputs feed a product-term allocator that improves flexibility in a couple of ways. First, it offers product-term steering, which allows completely arbitrary assignment of terms (i.e., O-16) to a particular macrocell. This capability goes well beyond the typical PAL which offers little (i.e., a couple of choices) or no steering, leading to the inevitable waste and sloppy fit.

A related efficiency enhancer is product-term sharing. With this feature, macrocells can use a common term and don't need to allocate multiple copies.

The I/O structure is a little different than that of PALs as well, splitting the registering (including D, T, latch, or combinatorial) and polarity functions from the output enable. This partitioning has a couple of benefits.

First, it deals with the fact that some applications are more I/O intensive than others. Thus, a given member of the Flash370 family features two variants (signified by even/odd part numbers, e.g., '374 and '375) that couple an I/O cell to every or every other macrocell, respectively. In the latter case, half the macrocells are described as buried.

Second, both the macro- and I/O-cells have feedback paths to the PIM, which addresses a chronic complaint of PAL users by allowing feedback without wasting a pin for an input.

One extremely useful feature of the '370 family is a simple timing model in which timing is both predictable and stable across design changes. Depending on the architectural specifics, CPLDs often exhibit myriad timing dependencies based on fanout, dedicated versus programmable pins, steering, sharing, and so on.

For instance, some CPLDs handle steering and sharing with the use of expanders that route product-term array outputs back through the array. This method improves flexibility and logic capacity, but at the expense of a second propagation delay. By contrast, the '370 timing is fixed for any configuration of steering, sharing, and I/O.

METALIABILITY

I don't want to encourage panic or paranoia, but in an era when you've got to sign a waiver to get a hot cup of coffee, product liability is clearly a concern.

Thus, it seems rather distressing to note that every electronic gizmo with a flip-flop (including PLDs) is shipped, even though (I can hear it now with suitably outraged tone, of course) the "negligent designer knew full well that it might fail." Sure, the MTBF may be a thousand years, but do you think that's going to stop a clever lawyer?

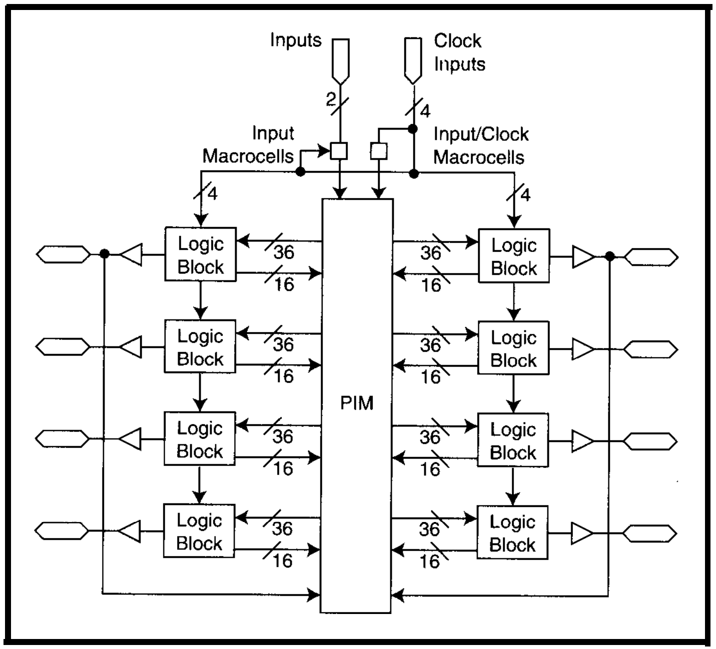

Figure 3—*The '370 family packs a number (Z-16 depending on mode/ number) of PAL-class logic blocks interconnected by a PIM (Programmable Interconnect Matrix).*

The problem is that an asynchronous input to a flip-flop is sure to eventually violate the setup and hold times. What happens then is the flop's output can enter an in-between 1 and 0 metastable state.

Most of the time, the result is no big deal (i.e., although it's called a walkout, it only means the flop's output eventually goes either high or low) within the specified clock-to-output delay. However, it is possible for the metastable state to persist long enough to produce an invalid output, which possibly propagates all kinds of bad vibes through the system.

There are a number of ways to attack metastability, both at the system design and device level. A designer might opt to use self-timed techniques (i.e., a flip-flop could provide a data-valid handshake signal only after the output resolves).

Alternatively, adding a second (or n^{th}) flip-flop pushes the probability of failure down accordingly. Recognizing that the situation is analogous to a ball poised on top of a hill, the steepness (i.e., gain-bandwidth product) can be increased to encourage the ball to quickly go one way or the other; a metastability hardening approach Cypress adopts that reduces, but doesn't eliminate, the problem.

However, by far the easiest (and thus most widely used) solution is simply to derate the system timing so the output is given more time to resolve. Notably, relatively small changes in timing can have big impact.

For instance, a PLD with a see-you-in-court metastability MTBF of a mere 10 ms at 41.6 MHz is granted immunity with a 20,000-year MTBF simply by dropping the clock rate to 33.3 MHz.

However, to have half a chance on the witness stand, this technique had better be backed up with some evidence that your timing choice (and thus accepted probability of failure) is reasonable (i.e., much less than other parts of the system, which shouldn't be hard, given the flakiness of software).

Cypress does the right thing and, in an aptly named "Are Your PLDs Metastable?" app note, quantifies the relevant characteristics of their parts, using a setup like that shown in Figure 5. In fact, they've also tested competitors' parts, but tactfully ask that you specifically request those results, rather than publishing them directly.

SYNTHETIC CHIPS

The Cypress PLDs are whizzy enough, but what really captured my imagination was their tool strategy.

PLD design tools have certainly come a long way since the invention of the PAL. At the very beginning, designers were actually encouraged to draw little xs on maze-like chip diagrams to designate unblown fuses, a technique optimistically deemed a "convenient shorthand" in the original MMI data book.

Subsequently, PLD designers moved onward and upward through PALASM

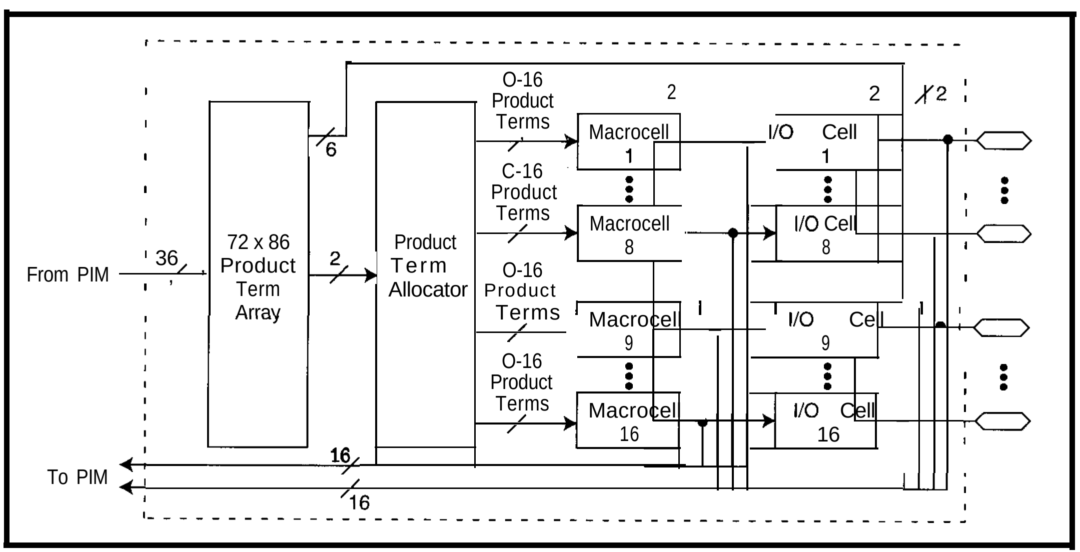

Figure 4—*Each '370 family logic block is arguably as powerful and feature rich as a standalone PAL's. I/O intensive family members (odd part numbers) combine an I/O cell with each macrocell and register intensive (even part numbers) with every other macrocell.*

and Boolean equations to more advanced tools like CUPL and ABEL with state-machine and schematic entry.

Cypress takes PLD development to new highs with a truly warped strategy. No, I don't mean crazy, touched, or twisted, but warped-as in their WARP family of VHDL (VHSIC Hardware Description Language) synthesis tools.

What is crazy is the fact the entry-level tool, WARP2, sells for an amazing $99! Crazy like a fox, that is, since the price is almost irresistible and sure to get Cypress phones (an 800 number that takes credit cards no less) ringing. Yes, WARP2 doesn't have all the features of high-priced synthesis tools, but it's by no means the feature-stripped, bait-and-switch deal you'd expect.

I first wrote about the then rather daring concept of logic synthesis five years ago ("VHDL-The End Of Hardware," *INK* 17). The VHDL language itself was defined at the behest of DOD (by IBM, TI, and others) explaining its somewhat Ada-like flavor. Since then, VHDL has achieved IEEE-standard (1076) status.

Now, hardware synthesis using VHDL (or the more C-like cousin Verilog] has become standard procedure for high-gate-count chip designers. There's really no choice in the matter, schematic entry having run out of gas as chips balloon beyond millions of transistors.

Besides dealing with the sheer complexity of modern chip design, synthesis yields other benefits including the ability to simulate a chip's operation long before silicon is available. This is extremely useful for everyone, including the chip designer (who can catch bugs and gotchas before it's too late), the system designer (who can integrate the chip's model into a system simulation), and the programmers [who can write code in the absence of real hardware).

Reflecting the title of the earlier article, Figure 6 shows hardware development using VHDL is conceptually little different than software development. Just substitute *compile* for *synthesize,* and you'll see what I mean.

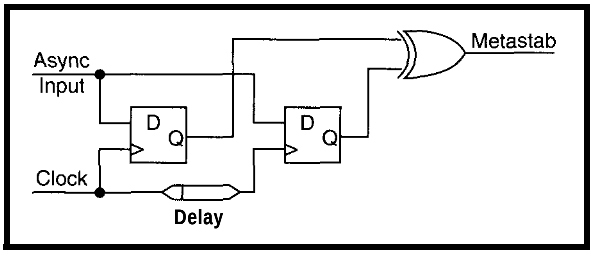

Figure *5-Cypress uses a test setup* **like this to evaluate and document the** *metastability* **characteristics** *of* **their** *PLDs.*

The figure also shows the various styles of programming supported by VHDL and WARP2, including behavioral, state table, structural, and Boolean. Let's take a look at an example of each in Listing **1.**

A behavioral description is easiest for the designer since, as the name implies, the designer only needs to simply define the behavior of the circuit with no concern for the implementation or timing of the circuits. For instance, the behavior of the circuit defined in Listing la could easily be specified in other ways, perhaps by replacing the `case` statement with a sequence of `if/then` clauses.

A state machine or lookup table is easy to define using the built-in `ttf` (truth-table function] operator. Listing lb describes a 7-segment LED decoder.

Structural descriptions differ from behavioral in that they explicitly dictate the underlying hardware. For example, the S-bit shift register in Listing lc unambiguously calls for a circuit using D-type flip-flops (`DFF`).

Finally, at the lowest (i.e., individual gate) level, VHDL supports the equivalent of PALASM-like Boolean equations (i.e., AND, OR, NOT, etc.), shown in the l-bit adder of Listing 1d.

As you can see from the examples, VHDL supports hierarchical design. A chip consists of multiple entities (a black box I/O definition much like the function prototypes in modern programming languages), each with an architecture (defining the behavior) which consists of one or more processes (code within a process executes sequentially with multiple processes executing in parallel).

A TIGHT FIT

Clearly, designers are best served working at the highest behavioral level since they need know nothing about the underlying hardware. Indeed, a big plus for synthesis is the ability to defer targeting a particular chip until the last second and to easily migrate between chips down the road.

Unfortunately, in the real world, true hardware independence is easier to talk about than achieve. Yes, you can synthesize any behavior to any chip, but total ignorance of the underlying hardware can be dangerous. The fact is that two behaviorally equivalent descriptions may synthesize to much different circuits depending on the fit between the functions called for by the synthesizer and those provided by the chip.

However, there are some generic programming recommendations at the behavioral level that can make for a more efficient (speed, size, power) design. A number of these are designed to work around a VHDL phenomenon known as *implicit memory.*

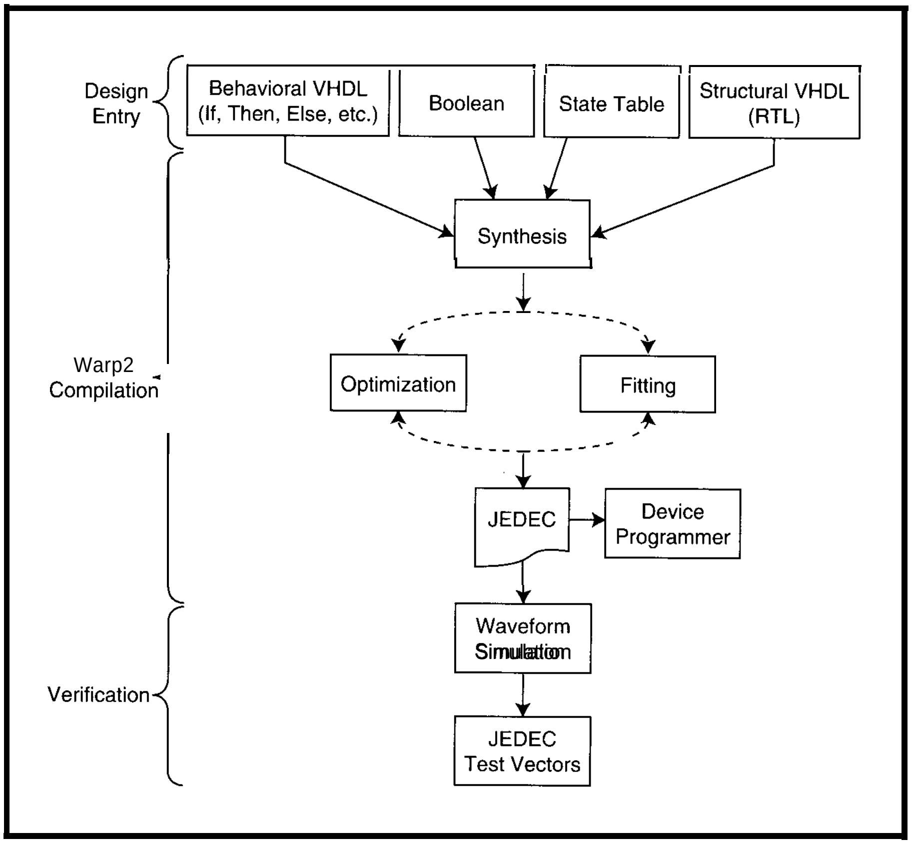

Figure 6—*Except* for fhe particulars of fhe language *used, methodologies* for hardware and *software* development *differ little.*

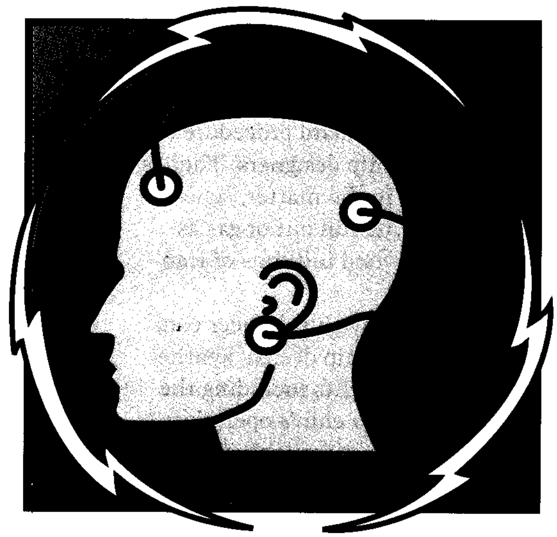

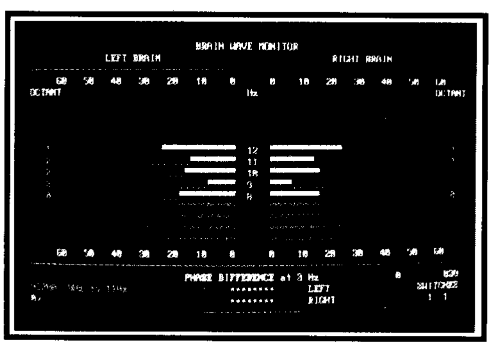

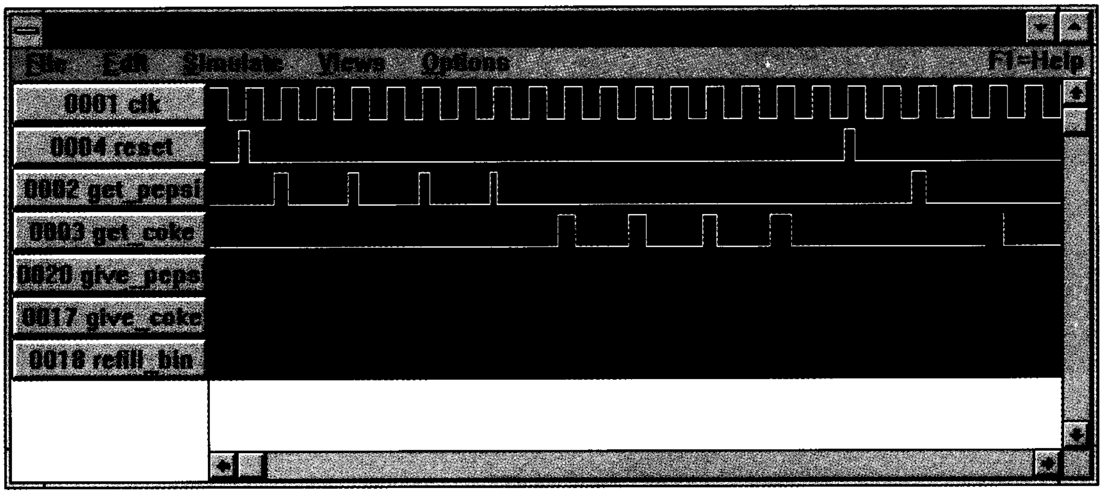

Photo 3—*Once the design is compiled, it's easy to simulate the circuit's operation. The included tutorial defines a simple soft drink machine that signals its need for refills once all the drinks (three of each) are sold.*

It turns out that VHDL doesn't rely on the explicit concept of a register. Instead, a register is inferred from a `wait` statement as shown in Figure 7a. The implicit memory issue arises when the statement is slightly modified with a conditional clause as in Figure 7b.

In this case, it's clear that `foo` takes the value `fred` if the condition is true (i.e., `jack` and `jill` both equal 1). But, what happens if the condition fails?

Some languages might decide the value of `foo` is then undefined or a don't care. However, VHDL presumes the value of `foo` should remain whatever it was and helps by synthesizing a latch. Needless to say, it doesn't take too much help like this to lead to excessive chip bloat.

However, simply completing the conditional expression as in Figure 7c solves the problem by synthesizing a multiplexer rather than a latch. It doesn't look much simpler until you realize the resulting circuit is easily optimized to that shown in Figure 7d.

Thus, getting rid of implicit memory is simply a matter of overtly specifying behavior rather than relying on assumptions by the compiler. Besides completing `if/then/else` chains with a final `else` clause as Figure 7c, other examples include terminating a case statement with a `when others` clause and initializing all outputs with a)

```
sync:process
  begin
    wait until clock = '1';
      foo <= fred;
  end;
```

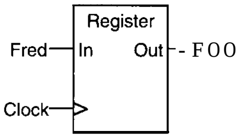

b)

```
sync:process
  begin
  wait until clock = '1';
    if (jack = '1' and jill = '1') then
      foo <= fred;
    end if:
  end process;
```

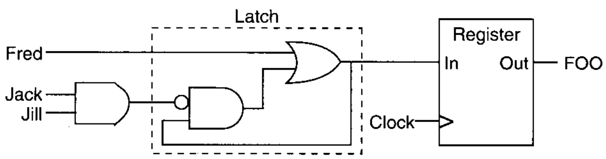

c)

```
sync:process
  begin
  wait until clock = '1';
    if (jack = '1' and jill = '1') then
      foo <= fred;
    else
      foo <= '0';
    end if:
  end process:
```

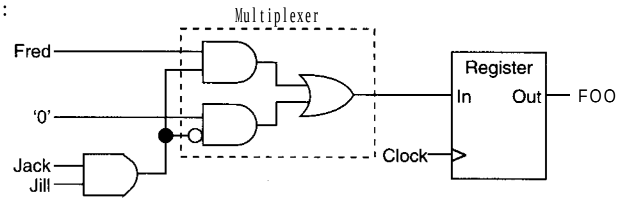

d)

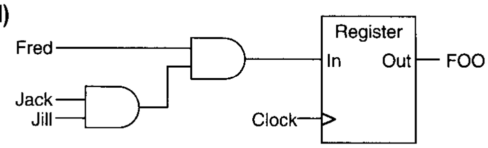

Figure 7—*Though* **a major goal** *of synthesis* **is to eliminate** *concerns about the underlying hardware, good programming practices yield more efficient circuits. Illustrated is an inferred register (a), open condition with latch (b), closed condition with multiplexer (c), and closed condition optimized (d).*

default values at the beginning of a process.

FULL SPEED AHEAD

WARP2, pictured in Photo 1, is available for Sun workstations and PCs running Windows. Reflecting the workstation heritage, the PC version calls for a rather beefy machine including **16** MB of memory [though some of it can be virtual), 55 MB of disk space, and a speedy CPU. The pain is soothed by the misery-loves-company fact that it takes a fully loaded PC to run a word processor these days, anyway.

Developing a chip with WARP2 is pretty much an iterative three-step enter, compile, simulate process. Design entry is completely textual, so you can use your favorite text editor.

Once the design is entered, it can be compiled without synthesis to quickly check for syntax errors. Once all the errors are eliminated, it's synthesized with the aid of a dialog box shown in Photo 2. Here, various real-world hardware constraints are chosen including specifying a particular chip, speed, or size optimization; manual or automatic pin assignment; and so on.

Listing 1—*VHDL language features cover fhe spectrum-from the high-level hardware independent behavioral /eve/ down to gate-level PALASM-like boolean constructs. These VHDL examples include a behavioral (a), fable (b), structural (c), and boolean (d) construct.*

a)

```
ENTITY sequence IS
  port (clk: in bit:
    s: inout bit);

ARCHITECTURE fsm OF sequence IS

SIGNAL count: INTEGER RANGE 0 TO 7;

BEGIN

PROCESS BEGIN
  WAIT UNTIL clk = '1';
    CASE count IS
    WHEN 0 |1| 2 | 3 =>
      s <= '1';
      count <= count + 1;
    WHEN 4 =>
      s <= '0';
      count <= count + 1;
    WHEN 5 =>
      s <= '1';
      count <= '0';
    WHEN others =>
      s <= '0';
      count <= '0';
    END CASE;
END PROCESS;
END FSM;
```

(continued)

Listing I-continued

b)

```
ENTITY SEG7 IS
  PORT (
    inputs: IN BIT_VECTOR (0 TO 3);
    outputs: OUT BIT-VECTOR (0 TO 6)
  );
END SEG7:

ARCHITECTURE mixed OF SEG7 IS

CONSTANT truthTable:
  x01_table (0 to 11, 0 to 10) := (
-- input &  output

   "0000" & "0111111",
   "0001" & "0000110",
   "0010" & "1011011",
   "0011" & "1001111",
   "0100" & "1100110",
   "0101" & "1101101",
   "0110" & "1111101",
   "0111" & "0000111",
   "1000" & "1111111",
   "1001" & "1101111",
   "101x" & "1111100", --creates E pattern
   "111x" & "1111100"
   );

BEGIN
    outputs <= ttf (truthTable, inputs);
END mixed:
```

c)

```
ENTITY shifter3 IS port (
    clk: IN BIT:
    x : IN BIT;
    q0: OUT BIT:
    ql: OUT BIT:
    q2: OUT BIT):
  END shifter3;

ARCHITECTURE struct OF shifter3 IS
  SIGNAL q0_temp, ql_temp, q2_temp: BIT:
  BEGIN
    dl : DFF PORT MAP (x, clk, q0_temp);
    d2: DFF PORT MAP (q0_temp, clk, ql_temp);
    d3 : DFF PORT MAP (ql_temp, clk, q2_temp);
    q0 <= q0_temp;
    ql <= ql_temp;
    q2 <= q2_temp;
  END struct;
```

d)

```
--entity declaration
ENTITY half-adder IS
  PORT (x, y : IN BIT;
    sum, carry : OUT BIT);
END half-adder:
--architecture body
ARCHITECTURE behave OF half-adder IS
BEGIN
  sum <= x XOR y;
  carry <= x AND y;
END behave;
```

Once synthesized, the Nova simulator is invoked to check out the design as shown in Photo 3 (the tutorial is for a simple soft-drink machine). It's easy to define inputs using a `clock` command to generate the CLK signal and clicking and dragging to specify the `get-inputs`.

Assuming everything looks good, a JEDEC fuse-map file is generated for use with a wide-variety of third-party programmers.

If you find WARP2 to your liking, consider moving up to WARP3 (at $4,995), which not surprisingly given the rather dramatic price difference, includes a number of big-ticket features.

First, it supports schematic entry, recognizing that that technique is still viable for relatively small designs and minimizing culture shock for traditionalists.

Second, it includes a fancy multi-windowed interface with source-level debugging.

Finally, a must-have reason to upgrade to WARP3 is that it also supports Cypress's pASIC38X line of antifuse-based FPGAs. FPGAs (like regular gate arrays) require much more complicated place and route than PLDs (the overnight compile option should give you a clue].

Tom Cantrell has been working on chip, board, and systems design and marketing in Silicon Valley for more than ten years. He may be reached at (510) 657-0264 or by fax at (510) 657-5441.

CONTACT

Cypress Semiconductor
3901 North First St.
San Jose, CA 95134
(408) 943-2600
Fax: (408) 943-2741
BBS: (408) 943-2954
http://www.cypress.com/
WARP orders: (800) 927-7843

I R S

428 Very Useful
429 Moderately Useful
430 Not Useful

CONNECTIME

conducted by Ken Davidson

The Circuit Cellar BBS
300/1200/2400/9600/14.4k bps
24 hours/7 days a week
(860) 871-1988—Four incoming lines
Internet E-mail: sysop@circellar.com

I like to try to pick out home-automation-related threads when the issue includes a **Home Automation & Building Control** *section, and I was successful at finding at least one. The first thread talks about adding a moisture-defection system to an HCS II (or any other controller) to hand/e the task of watering the lawn.*

Next, we look at sensing the opacity of a column of smoke. I suppose those people frying to automate their home wood stove could fake things to the extreme with this one.

Finally, we talk about some options for soft-starting an AC motor in an effort to extend ifs useful life.

HCS II lawn system

Msg#: 4942
From: Mike Cavanaugh To: All Users

I am busy thinking of more tasks for my HCS II to handle. With my large yard, it is tough to spend the weekend watering the lawn. So, my next project may be to install an automatic lawn sprinkling system.

Some of my concerns:

- How do I detect if the lawn needs to be watered?
- *Popular Electronics* showed a "moisture detector" based on burying two probes about 2″ under the ground and then monitoring the resistance. Has someone done this or have they got a better idea?
- Do the store-bought moisture detectors work OK-for example, the RainMatic RainSensor (p/n HCC-RMRS), available from Home Control Concepts for around $33? Is there a better or cheaper unit available?
- Do the water on/off solenoids tend to stick or burn out frequently?
- Does anyone have good experience with a certain type of solenoid?
- Are there other items that I should be concerned with?

Msg#:10455
From: Pellervo Kaskinen To: Mike Cavanaugh

The simplest soil moisture sensor is likely to be the one based on a battery principle. Get some zinc and some copper (for example]. Make two small stakes or nails out of them and drive them in the soil, close to each other, but not touching. Connect a wire to each electrode. At the other end of those wires, connect a load resistor. You probably want to run about 50 µA through it, and as the zinc to copper voltage is around 1 V, it would mean 20 kΩ.

The higher the moisture level, the more current the chemical cell can support. That means, your voltage across the load resistor increases towards the stated 50 µA, but also saturates because of the limiting voltage of the process. That means you want to experiment with the resistance value based on the distance of the two electrodes from each other and the "salt" content (i.e., the fertilizers, acids, or lime) of the soil, in addition to the moisture.

Basically, think of the cell as a metal-pair-dependent voltage source VO connected in series with an electrolyte-dependent ("salty" soil water) source resistance RO. Then the load resistor R1 sees VO*R1/(RO+R1).

This process is used in some commercial houseplant water-need indicators. In fact, you might pay a visit to a nearby florist. The commercial unit just uses a simple microampere meter (with color bands instead of a graduation) as the load. Unlike the principle I outlined with the two nails, the commercial unit is built as a single coaxial element. It resembles a phone plug, with the tip region being one metal and the shank region the other metal. There is a built-in insulator between the two sections. This way you get repeatable results from unit to unit.

I don't have experience with the type of solenoid valves used in the lawn watering systems, but ordinary industrial solenoid valves suffer mainly from one problem: clogging of the pilot orifice. If your water has plenty of minerals and rust, this is likely to happen on any valve.

Another thing happening to solenoid valves after some years is the sticking of the plunger inside the solenoid. This normally prevents the valve from closing. But, it is easy to clean in most cases.

Msg#:12140
From: Mike Cavanaugh To: Pellervo Kaskinen

Interesting idea. I will have to try it out and see how stable the ground proves to be in relation to the chemicals present and how or if it changes much. I am still inclined to bury two pieces of metal and then use an op-amp to measure the resistance and as the moisture varies, the resistance of the op-amp feedback circuit varies, causing the op-amp to output a voltage which represents the moisture content of the ground. Still in the thinking stage. Thanks.

Msg#:14159
From: James Meyer To: Mike Cavanaugh

If you do use two metal electrodes for a resistance measurement, you should think about feeding them with an AC voltage instead of DC.

With DC, you'll likely get some electrolytic action that causes the resistance to drift quite a bit. The electrodes will "charge" just like a battery. With AC, that can't happen. It would be only a bit more complicated to use AC.

Msg#:14453
From: Pellervo Kaskinen To: Mike Cavanaugh

If you want to measure the resistance, you *have* to use an AC circuit. The same effect that makes the "battery" will foul any ordinary electrodes in a DC resistance-measuring scheme in a few seconds, or minutes at most.

Msg#:14459
From: Mike Cavanaugh To: Pellervo Kaskinen

One quick thought: what if I were to use stainless electrodes? I am thinking that any buildup would be small. I like DC because of the safety factor.

Msg#:19075
From: Pellervo Kaskinen To: Mike Cavanaugh

Stainless steel does not prevent the fouling I mentioned. Even platinum does not. It is a build-up, not a corrosion issue.

When we talk about AC, it is not a safety risk. What we need, is simple alternating polarity that you would create with a multivibrator. Nobody suggests you would use the household 115 V for the measurement.

Still, my suggestion about the flower pot moisture sensors should be the simplest and potentially cheapest way.

Msg#:19710
From: Mike Cavanaugh To: Pellervo Kaskinen

How high of an AC voltage is required? Maybe a small transformer (12 VAC or 24 VAC) would do the job? Or is this a "try and see" type of thing. I will look for the "flower pot" things. Thanks.

Msg#:24061
From: Pellervo Kaskinen To: Mike Cavanaugh

I probably would try to do it on a 6.3-V transformer, if I would choose to go the AC bridge measurement route. If and if ;-)

Measuring smoke outdoors

Msg#: 4295
From: Richard Cooke To: All Users

I have an application that uses the PIC 16C73 and needs to measure the opacity of a column of smoke outdoors. In the lab it is an easy task, but I need to make it immune to ambient light (60 Hz or sunlight). I don't want to reinvent the wheel, so I hope some kind soul here might have a few words of wisdom. If I pulse the LED what frequency is best? What is the recommended receiver circuit for pulsed input?

Msg#: 4296
From: Russ Reiss To: Richard Cooke

A technique that has provided good ambient light immunity in IR touch panels in the past works like this: The output of the light sensor is fed to a track-and-hold circuit. Just prior to turning on the LED for opacity measurement, you hold the last ambient light reading. Then you pulse the LED and subtract the held (last) ambient reading from the sensor's response. This presumes that the ambient light (sunlight or 60 Hz or 120 Hz or whatever) doesn't change much during the brief sampling time. Might work for you too.

Msg#: 4303
From: Richard Cooke To: Russ Reiss

What about doing it all with the PIC? I could read the ambient light and then turn on the LED and read again. Subtracting the two readings should work.

Msg#: 4438
From: Pellervo Kaskinen To: Richard Cooke

Even though Russ may already have pointed this out, I'll try to elaborate. If your signal is more or less swamped by broad daylight, then there is little chance that the 256 discrete states resolution of the PIC analog-to-digital front end can resolve any meaningful smoke density data.

You have to do more at the analog front end. You have to make the DC operating point of your photosensor such that it is not completely saturated with the daylight. Then you have to AC couple it and bandpass filter it. The passband frequency must be the same that you use for modulating the transmitter.

After the bandpass filter, you probably need a rectifier circuit-I think a PIC ADC is limited to single-polarity signals. You probably also want to use a little smoothing after the rectifier rather than letting the PIC to do all the averaging. I mean, what is a single capacitor among friends?

Msg#: 4390
From: George Novacek To: Richard Cooke

I don't think pulsing is going to do you much good at low (below 10%) obscuration level, which is probably what you need. It's been a few years since I worked in this area, but in general, there were methods to minimize the ambient light effects at low obscuration levels.

You could use pure, narrow bandwidth color, which is not present in the ambient light. Near-infrared is often workable, but selective filters are very expensive and you never get 100% selectivity anyway.

More practical method lies in building light baffles, such as in photoelectric smoke detectors, which allow free passage of air but stop ambient light. In addition, measurement of light scattering as opposed to light obscuration might give you better results. You must also remember that the light sources age, as do photodetectors. They also get dirty. So you will have to consider calibrating at regular intervals.

I don't know of any books on the subject. To see what others have done, buy a photoelectric detector and study it. Also, go to a patent office and do a search of prior art. Pull the old patents and you'll be surprised how much you will learn.

Msg#: 4416
From: George Novacek To: Richard Cooke

One more thought on the subject: In the days before the microprocessor, we were always taught that the best receiver is a good antenna. Since the microprocessor appeared on the scene, just about one out of two product designs I have seen have been outright atrocious. Engineers seem to have adopted a cavalier attitude, assuming that the *almighty* microprocessor is a panacea and its *intelligence* will fix every inadequacy. It won't.

You can't bleed stone. If your signal is flaky, you are off to a bad start. Before you even begin considering fancy ways of extracting useful signal from the ocean of garbage, make sure it is as good and clean as realistically possible. In your case, I would look at proper design of an efficient optical system first. Then, and only then, start looking at fancy signal processing methods.

Soft starting a small AC motor

Msg#:21028
From: Jim Fletcher To: All Users

I am trying to give a small AC motor (<1/10 HP) a soft start. I believe if I could in some way ramp up to 110 VAC instead of applying full voltage to the motor, it will increase the life of the motor. The motor is constantly being turned on and off. If anyone has any ideas or circuits in mind, please let me know. Thanks!

Msg#:34759
From: Pellervo Kaskinen To: Jim Fletcher

There are several concepts for the soft starting of AC motors. Which one is the most appropriate depends on the exact application and the size of the motor *relative* to the load. It also depends on the type of the motor.

Let's list the motor types first. You can have a universal motor, which actually is a DC motor optimized with low inductance so it can work while the polarity of the power line alternates.

Then there are the capacitor-start (and capacitor-run, most likely) motors. They are basically like 3-phase motors, with the capacitor passing phase-shifted current to one of the windings.

And last, there are the inductively phase-shifted motors, where part of the poles is separated and surrounded with heavy copper loop. This loop opposes any change of the field within and thereby delays the magnetic %-pole. Consequently, you get a rotating field for the motor to follow.

Of these three principles, only the capacitor-run version can be reversed in the rotation direction, without any really elaborate systems like clutches and gearboxes. Another characterization is that the shaded pole [shorting loop) type is normally seen in the smallest motors, up to about 50 or 100 W only. The shorting loops cause poor efficiency, which becomes an issue at higher powers. The other two types can be obtained for higher powers, with the universal motor commanding the domestic vacuum cleaner market and the capacitor type being used in the air conditioning and the refrigeration devices.

Now, assuming that your 1/10-HP (75-W) motor is a capacitor-run motor that also starts on its run capacitor rather than a separate start capacitor, we have a few choices for the soft start.

1. You could have a completely independent soft-start element on the shaft. Such elements are available in powder-filled, hydraulic, and eddy-current-dependent clutches. I think you do not have these in your wish list.

2. You can indeed ramp the voltage up gradually. Conceptually, the simplest thing would be to apply a variable transformer to supply your motor with the ramping voltage. But then, what would adjust the variac? Probably you do not want another motor to do that. Sigh.. .

3. You could ramp the supply up with a power transistor mounted from the positive terminal of a diode bridge to the negative terminal, while the AC comes to one side AC terminal and the other AC terminal of the bridge goes to the motor. The other side of the motor completes the loop to the power source. Controlling the base of the transistor you can do whatever ramping you want. The two difficulties with this approach are the heat dissipation and the need for isolation to the transistor base. It runs at the power line potentials.

4. A crude (or simple and elegant, if you agree) solution is to put in a series inductance (a choke) of an appropriate value. With the high current during initial power-on, the choke drops much of the voltage, so the motor sees a reduced voltage and has a lower torque. Thereby it spins up slower. But once it does spin up, the current is reduced, and the choke drops less of the supply, leaving more to the mo-

tor for continued acceleration. The choke, not being resistive in nature, does not heat up too much, either.

Now, there are a few conditions. The motor has to have enough torque to overcome any static friction of the load, even at the reduced voltage. The load cannot be of a fixed-torque type. Rather, it has to be something like what a fan produces, with increased torque requirement with increased speed. If it is fixed, it has to be very small indeed.

An AC motor of the capacitor-run type has a poor starting torque. A DC motor or a universal motor has a starting torque equal or higher than the running torque, but the AC motor has a starting torque of roughly a fifth to a third of the running torque.

The torque developed by an AC motor is proportional to the *square* of the voltage. If you cut the supply voltage to half, your starting torque is reduced to a quarter. The current draw is at the same time cut to about a half.

Consequently, you probably want to try an impedance that drops the starting current to about half as your first shot on the problem of dimensioning the beast. Remember that the selection ultimately depends on the load as well as the motor.

Your 75-W motor is likely to draw about 1 A of current at run time and 3–10 times that during the first moment of a start, given no additional impedance. This forms a basis for our estimate. You would want the choke to drop about 60 V at about 3 A (half of 6 A that I assume is within the 3-10 range, for lack of better info). That means its impedance should be 20 Ω at 60 Hz, or an inductance of 53 mH.

5. Triac-based systems like the lamp dimmers may be the more acceptable type for you. They are readily available and you could tinker with them more easily than with the choke value estimated above. They also would produce an independent ramping rather than the motor-speed-dependent one that is more prone to leaving the motor in a stalled condition if the load ever increases from the value originally used for the trimming.

My suggestion is that you would buy a cheap light dimmer and study the situation with it. The dimmer may not be immediately suitable for the automatic ramping because it is AC based, but we'll come to that issue shortly.

You probably need some sort of noise filtering because of the sharp rise of the voltage every time the triac fires, although the motor is inductive enough to take care of most of the issue. This filtering would also include the RC snubber values over the triac.

To make the circuit work in a ramped fashion, we generally want to make it work in a DC environment. Return to the idea of the power transistor in a diode bridge. Your light dimmer just might work in the same fashion. Then, you would route the feed to the dimmer potentiometer through a resistor/capacitor stage that having a sufficiently large capacitor allows only a slow build-up of the firing angle. The resulting voltage at the load would ramp up slowly.

I must emphasize that I have not tried any light dimmer lately and there might be surprises in the way they are implemented. But there also are surprises and hassles with all the other approaches that I have outlined. Take your pick!

We invite you to call the Circuit Cellar BBS and exchange messages and files with other Circuit Cellar readers. It is available 24 hours a day and may be reached at (860) 871-1988. Set your modem for 8 data bits, 1 stop bit, no parity, and 300, 1200, 2400, 9600, or 14.4k bps.

ARTICLE SOFTWARE

Software for the articles in this and past issues of *Circuit Cellar INK* may be downloaded from the Circuit Cellar BBS free of charge. It is also available on the Internet at **ftp://ftp.circellar.com/pub/circellar/**. Web users should point their browser at **http://www.circellar.com/**. For those with just E-mail access, send a message to **info@circellar.com** to find out how to request files through E-mail.

For those unable to download files, the software is also available on one 360 KB IBM PC-format disk for only $12. To order Software on Disk, send check or money order to: Circuit Cellar INK, Software On Disk, P.O. Box 772, Vernon, CT 06066, or use your Visa or Mastercard and call (860) 8752199. Be sure to specify the issue number of each disk you order. Please add $3 for shipping outside the U.S.

IRS

431 Very Useful 432 Moderately Useful 433 Not Useful

PRIORITY INTERRUPT

The Old Curmudgeon

Perhaps because I emphasize practical application over fantasy technology and overblown specifications so much, many people might think I'm just an old curmudgeon. Despite the fact that I'm on the leading edge of the boomer bubble, I don't lament too often about the past, and I indeed appreciate the technological advances of our present era. I'm not stuck thinking that a four-function calculator is the most computing power I'll ever need, but neither have I become a "Cache-Crazed Penthead" who trades in his computer each time the processor changes or clock speeds increase 10 MHz.

Still, we have to be wary of the craziness, All the excitement continues to promote announcements of vaporware because companies are forced to vie for positions in barely emerging markets. The bad news is that when certain big companies announce products-real or not-smaller ones often don't bother to compete. The end result is that we get to wait while the big guys take their time making it real and often end up with few other choices, The worse consequence occurs when it turns out that the vapor had more vitality than the delivered product.

Like many of you, I went out and bought Windows 95 (there were no lines where I bought it) to upgrade my Gateway Pentium. It seemed funny to upgrade a three-month old machine, but you know, there was all that hype.

Personally, I think the last good software I've seen was written for 64-KB Z80-based machines about 15 years ago. Back then, you had an upper limit on available memory, and if you wanted to add more features you programmed it better, not bigger. Because of the concise assembly routines and lack of leviathan graphic manipulations, execution speeds were extremely fast, even with 8-bit processors. I still think SuperCalc ran faster on my Digital Group Z-80 than these new spreadsheets do on a Pentium (doing a simple function, of course). These days, every time I get a new release of software, it's 3-5 MB larger and needs the next generation processor just to keep performance even with the old version.

As I passed Jeff's office this morning, he was still swearing after two days of trying to "upgrade" his machine to Windows 95. Do you call it a "downgrade" when the result is worse? :-)

Hardware isn't much better this year. Improvements have become much more incremental. While we have seen the rapid introduction of the 80486, Pentium, and Pentium Pro (except in a few cases where the Pentium Pro uses wider data paths), they are all still basically 32-bit processors, albeit each with clock rates faster than its predecessor. If the main reason to need a faster system is to run the next generation software at the same performance level (I already admitted I have a rather warped view on this subject), wouldn't it be better to take a breather and pay attention to getting software technology to truly use one processor level before jumping to the next?

Finally, there's the Internet. Some publicists are calling it the citizen's band radio of the '90s. Great! For any of you who experienced the citizen's band radio the last time, you know what a zoo it's going to become. Time to switch hobbies, lease a private gateway, or just sit this one out-who knows? Like everything else in this game, people's expectations are surely going to hype it out of the practical.

As for myself? When the screams from Jeff's office subside, I'll remind myself that "If it ain't broke, don't fix it" and find another reason to postpone the inevitable "upgrade."

Steve

steve.ciarcia@circellar.com