

CIRCUIT CELLAR[®] INK[®]

THE COMPUTER APPLICATIONS JOURNAL

#78 JANUARY 1997


EMBEDDED APPLICATIONS

Add a 68HC11 to
an RC Vehicle

Microchip's
PIC14000 Battery
Management Chip

Applying GPS


Ciarcia's HCS II
Weather Station



\$3.95 U.S.
\$4.95 Canada

TASK MANAGER

New Toys



ately, I've been pondering what toys I'd like to get in the new year. I'm always on the lookout for the latest gadgets that might make my life easier. Of course, lusting after every little gizmo on display at CES doesn't do much good if you can't buy it on the open market, so my desires have to be tempered with a touch of reality.

As it turns out, the two devices on the top of my list are perfect examples of embedded applications. They show just how deep computers are permeating every aspect of our lives.

I've been using a Sharp Wizard PIM for the last few years to keep track of my schedule, telephone numbers, and miscellaneous notes. However, it's somewhat big and bulky, making it less than convenient to carry everywhere I go.

So, topping my list is the Timex Data Link watch. I'm sure you've seen these beauties. You enter your phone numbers, appointments, and to-do list on your PC, hold the watch up to the computer screen, and the information beams over from the PC to the watch via flashing bars of light on the screen. Always on hand (literally), it's the perfect PIM when you don't need to change information on the fly.

I want my next toy to help me keep track of my thoughts. I think best in the shower and while driving (though not at the same time). In both cases, it's not convenient to grab a pencil and paper to jot down notes. The alternative is a bulky microcassette recorder that is worse to carry around than the Wizard.

Enter solid-state recording chips. It's now possible to make a small, lightweight voice recorder for very little money that's perfect for those short little notes. Between the recorder and the watch, I hope to add just a bit more organization to my hectic life.

You need only flip a few pages here to see more examples of embedded applications. First, Bill Payne describes how to use some simple chips to make a low-power, solid-state DC-motor-control circuit. And, then, what would an embedded applications issue be without a mobile project? In our next feature, Dean Hoyt adds some smarts to an RC car so it can record a sequence of movements and then replay them exactly.

Next, Fred Eady crosses over from the *Embedded PC* section to look at a special-purpose PIC with integrated analog functions that makes a perfect battery-management chip. Brian Millier adds a normally expensive piece of test gear to his home lab by designing it himself. And, Do-While Jones completes his discussion of GPS by exploring some implementation issues.

In our columns, Steve expands on the HCS II's already impressive network capabilities with a new network module. Jeff displays his knowledge of LCDs with a theoretical and practical discussion. Finally, Tom reports back on what heated up the latest Hot Chips show.

editor@circellar.com

CIRCUIT CELLAR[®]

THE COMPUTER APPLICATIONS JOURNAL

FOUNDER/EDITORIAL DIRECTOR
Steve Ciarcia

PUBLISHER
Daniel Rodrigues

EDITOR-IN-CHIEF
Ken Davidson

PUBLISHER'S ASSISTANT
Sue Hodge

MANAGING EDITOR
Janice Marinelli

CIRCULATION MANAGER
Rose Mansella

TECHNICAL EDITOR
Elizabeth Laurençot

CIRCULATION ASSISTANT
Barbara Maleski

ENGINEERING STAFF
Jeff Bachiochi

CIRCULATION CONSULTANT
John S. Treworgy

WEST COAST EDITOR
Tom Cantrell

BUSINESS MANAGER
Jeannette Walters

CONTRIBUTING EDITORS
Rick Lehrbaum
Fred Eady

ADVERTISING COORDINATOR
Dan Gorsky

NEW PRODUCTS EDITOR
Harv Weiner

ART DIRECTOR
Lisa Ferry

PRODUCTION STAFF
John Gorsky
James Soussouis

CONTRIBUTORS:
Jon Elson
Tim McDonough
Frank Kuechmann
Pellervo Kaskinen

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. Periodical rates 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 (660) 875-2199.

HAJAR ASSOCIATES NATIONAL ADVERTISING REPRESENTATIVES

NORTHEAST & MID-ATLANTIC
Barbara Best
(908) 741-7744
Fax: (908) 741-6823
B.Best-Hajar@worldnet.att.net

MIDWEST & SOUTHEAST
Christa Collins
(954) 966-3939
Fax: (954) 985-8457
HajarChrista@worldnet.net

WEST COAST
Barbara Jones
& Shelley Rainey
(714) 540-3554
Fax: (714) 540-7103
shelley.hajar@worldnet.att.net

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/>

All programs and schematics in Circuit Cellar INK[®] have been carefully reviewed to ensure their performance is in accordance with the specifications described, and programs are posted on the Circuit Cellar BBS for electronic transfer by subscribers.

Circuit Cellar INK[®] makes no warranties and assumes no responsibility or liability of any kind for errors in these programs or schematics or for the consequences of any such errors. Furthermore, because of possible variation in the quality and condition of materials and workmanship of reader-assembled projects, Circuit Cellar INK[®] disclaims any responsibility for the safe and proper function of reader-assembled projects based upon or from plans, descriptions, or information published in Circuit Cellar INK[®].

Entire contents copyright © 1997 by Circuit Cellar Incorporated. All rights reserved. Circuit Cellar INK is a registered trademark of Circuit Cellar Inc. Reproduction of this publication in whole or in part without written consent from Circuit Cellar Inc. is prohibited.

- 1 4** Low-Power, Solid-State DC Motor Control
by Bill Payne
- 2 0** A Radio-Controlled Car and a **68HC11** Microprocessor
by Dean Hoyt
- 2 6** Using Microchip's **PIC14000** Battery-Management Chip
by Fred Eady
- 3 4** Pulse'05 Pulse Generator
by Brian Millier
- 4 2** The Global Positioning System
Part 2: Implementation Problems and Solutions
by Do-While Jones
- 5 0** **MicroSeries**
Applying the HCS II
Part 2: A New Approach to High-Performance System Expansion
Steve Ciarcia
- 6 8** From the Bench
Liquid Crystal Displays
First You See 'Em, Then You Don't
Jeff Bachiochi
- 7 4** Silicon Update
Haute Chips
Tom Cantrell

INSIDE ISSUE 78

- 2** Task Manager
Ken Davidson
New Toys
- 6** Reader I/O
Letters to the Editor
- 8** New Product News
edited by Harv Weiner

- ConnectTime**
Excerpts from
the Circuit Cellar BBS
conducted by
Ken Davidson
- Priority Interrupt
Steve Ciarcia
Hardly a High-Definition
Opinion
- Advertiser's Index

96

65

READER I/O

STAMP OF APPROVAL

As I was browsing through the October 28, 1996,

Gassée and the BeBox. Aside from it being an interesting article, I couldn't help but notice that the photo of Mr. Gassée showed a copy of **INK** prominently displayed on his desk.

Clearly, this is a man to be taken seriously. Had he been reading **Business Week**, I wouldn't have given much consideration to the future of his company.

Bob Lawson
robert.lawson@epamail.epa.gov

ARE YOU HAVING FUN YET?

I want to stress the importance of humor in technical journals. There's one phrase from an **INK** article that still makes me smile. "Two pins are used for power and ground (no one's thought of a way around that...yet)."

Some would snip that sort of thing in vague homage to the gods of professionalism, but that's wrong! Yes, often what we do is tedious, dry, and boring. But if we're not having fun, we've lost our creativity.

Keep the human being in engineering. We need both the paragraph on a groaning debug session and that sarcastic industry insight.

Bob Campbell
sfdstrct@netwizards.net

THAT TECHNICAL EDGE

Steve's "Bits and Bytes" (**INK 76**) made me think back to a much different time.

It was the early '70s, and I'd just received **BYTE's** first issue. I pored over the articles, gaining that much-needed practical knowledge that a newly graduated engineer requires at the beginning of a career. I enjoyed your hardware approach to things in those early days.

As the years went on and PCs became an important engineering tool, **BYTE** was always on top with good technical hardware and software articles. But, something happened. The good technical articles became reviews of the 12 fastest printers, the best operating system—articles that appealed to computer techno jerks, not engineers.

Then, along came **INK**, and I was back in technical heaven. I now manage an engineering department and do much more managing than engineering, but I still wait

for my issue each month. I pass my **INK** issues on to all my young engineers to give them that same experience.

Thanks, Steve. Keep us on that technical edge, and please don't give in to stockholder demands and destroy something good.

Bob VanNaarden
rvannaar@wirenet.com

"HEY! WATCH WHERE YOU'RE GOING!"

Steve's "The Radical Fringe" (**INK 75**) reminds me of the scrolling LED display I built from scratch for a college project. The prof wanted a vehicular display to flash messages to people driving with less than optimum courtesy. I don't know if he ever had the guts to use it, but it was fun to build. I still remember the astonishment on his face when I demonstrated it. It was obvious he'd never built a working microprocessor-based device!

And, please keep ragging on the high coefficient of ridiculousness and lack of common sense and usability of the DOS and Windows software universe. I have no tolerance for such nonthinking inefficiency.

Hank Wallace
hwallace@intrlink.com

Contacting Circuit Cellar

We at Circuit Cellar **INK** encourage communication between our readers and our staff, so we 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 offices at (860) 875-2199.

Fax: All faxes may be sent to (860) 871-0411.

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.com**. 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/**.


NEW PRODUCT NEWS

Edited by Harv Weiner

VIDEO CHARACTER GENERATOR

BOB-I enables small microcontroller-based systems to display substantial amounts of text either as an overlay on composite video from standard **RS-170A** NTSC cameras or by generating a complete video signal onboard. The output signal from **BOB-I** can be distributed using low-cost CCTV technology for multiple (or giant) display applications (e.g., electronic signs). In many cases, it connects to an SBC or microcontroller with only one spare I/O pin (four max.). Its video input circuit is differential, so it can reject low-frequency common-mode noise and prevent ground-loop problems.

BOB-I features a character format of 12 x 18 pixels with 240 maximum displayable (10 rows, 24 columns). Character-display options include seven character colors, eight outline/background modes, blink, and size increase up to



four times. A character ROM includes 128 patterns with some character graphics. A 16-bit synchronous serial data path is used with a maximum data rate of 2.5 Mbps (400-ns min. bit-clock cycle time). The board measures 4.00" x 2.20" x 0.70" and is powered by +12 VDC nominal at 80 mA max. Connections to the board are through 0.025" square post headers.

BOB-I sells for \$186 in quantities of three or fewer. Discounts are available for higher quantities. A **BOB-I Developer's Package**, available for \$200, includes demo software (with C source code) that exercises the board through a PC parallel port.

Decade Engineering
2302 5th St. NE
Salem, OR 97303-6832
(503) 363-5143
Fax: (503) 399-9747

#500

I/O CONTROLLER

ISA-Anvil is a flexible development platform for ISA-bus systems. It features a 3200-gate FPGA mounted on an ISA-compatible circuit board especially designed for hardware prototyping. The unit includes everything needed to develop specialized PC-based

peripherals using reconfigurable FPGA technology.

The 132-pin FPGA connects to the ISA-bus address, data, interrupt, and DMA control lines. So, the board is easily configured as an I/O or memory-mapped peripheral, as well as an interrupt or DMA control-

ler. In addition to the system interface, 60 FPGA pins are available for programmable I/O functions. These 60 PCI-compliant I/O lines, which can sink 24 mA or source 16 mA, are connected to ribbon cable headers.

ISA-Anvil uses the Altera **EPX880** FPGA, which operates at a toggle frequency of up to 80 MHz with gate delays less than 10 ns. The SRAM-based **EPX880** architecture features ten 22V10-compatible logic array blocks. Any array block can also be configured as 1280 bits of static RAM. Power consumption of the gate array is under 2.5 mA per MHz.


ISA-Anvil includes Altera's **PLD-Shell** development software, which has a **PALASM2**-compatible com-

piler, estimator, and functional simulator. Error files specify syntax or device fitting errors, and waveform files display simulation results. Compiled JEDEC files load directly into the FPGA through the chip's JTAG port.

ISA-Anvil sells for \$199.

Techniprise, Inc.
459 Highgate Ave.
Worthington, OH 43085
(614) 431-5610
Fax: (614) 431-5617

#501



NEW PRODUCT NEWS

SRAM Module

A new line of surface-mount NV SRAM modules featuring a built-in battery monitor and CPU reset as well as reduced power consumption has been announced by Dallas Semiconductor. Series devices include the DS1330 (32K x 8), DS1345 (128K x 8), and DS1350 (512K x 8).

Each DS13xx module contains a sophisticated battery monitor that periodically checks the voltage of the module's internal lithium battery. When the battery voltage is less than 2.6 V, the battery warning pin is activated to signal the need for module replacement. DS13xx products offer 10 years of data retention at 25°C.


These devices are designed to be low power. Although existing NV SRAMs have standby currents in the range of 3-10 mA, the new modules have typical standby currents of only 50 μ A, so they're suitable for portable and hand-held systems.

The devices are available in an exclusive surface-mount Low Profile Module (LPM) package. This 34-pin package snaps directly into a standard 68-pin PLCC socket, providing a surface-mount solution that protects the internal lithium battery from the high temperatures of reflow soldering. Users simply mount a PLCC socket during soldering and then snap an LPM module into the socket.

The DS1330 is priced at \$16.12 apiece in quantities of 1000.

Dallas Semiconductor
4401 S. Beltwood Pkwy.
Dallas, TX 75244-3219
(972) 371-4448
Fax: (972) 371-3715

#502



ZERO EXTERNAL-COMPONENT MICROCONTROLLER

The COP8SAx family of One-Time-Programmable (OTP) devices integrates many features found externally on most microcontrollers, creating a true complete system on a chip for 8-bit applications. The family saves PCB space as well as manufacturing and passive-component costs, while increasing reliability.


The family features a range of devices based on an industry-standard core architecture. RISC-like performance is achieved via the 1-us instruction cycle time and the fact that over 77% of the instructions are single byte and single cycle. Various combinations of RAM, EPROM, user storage space, and package size are available.

The COP8SAx offers full integration with R/C oscillator, power-on-reset, pull-up resistors, Schmitt triggers, and protection diodes—all internal to the chip. The chip also features a 16-bit timer/clock, two low-power modes, wakeup on a pin change, watchdog timer/clock monitor, direct LED drive, enhanced serial interface, and the ability to use more software in a smaller space. Low 2.7-5.5 operating voltage enables designers to implement intelligence in low-power applications.

User-friendly development tools range from a free software simulator to a low-cost Evaluation and Programming Unit (\$99), a midrange Debug Software Module (\$699), or a full-featured iceMASTER 400 (\$2990) for in-circuit emulation. Software development tools (e.g., the C Compiler (\$795) and free COP8 Assembler/Linker/Librarian) take advantage of the chip's new features.

National Semiconductor Corp.
2900 Semiconductor Dr. • Santa Clara, CA 95052-8090
(408) 721-8299 • Fax: (408) 735-7978 • <http://www.national.com/>

#503



NEW PRODUCT NEWS

EMBEDDED-PC MODULES

The Talix **i386EX** is smaller than a credit card, but it expands into a highly compact and powerful '386-based PC-compatible controller. Optional modules provide links to Ethernet and CAN networks, and access to PCMCIA-based storage and communication peripherals.

The **i386EX** contains a 25-MHz '386 processor, 1-MB RAM, and OS-MB flash memory on a board measuring just 2.0" x 3.1". Also onboard are two serial channels, a watchdog timer, three separate timer output channels, interrupt control, parallel I/O lines, and expansion connectors for extra interfaces. Two power-down stages ensure optimum power management even though normal power consumption is less than 500 mA at 5 V.

The **i386EX** is available with Datalight's MS-DOS 6.22-compatible ROM-DOS, simplifying software development for PC-compatible target systems. Prototyping embedded applications is further simplified with the **i386** development board. Acting as a motherboard for the **i386**, this board provides a prototyping area for application-specific circuitry and status LEDs for debugging. Onboard are ISA and PC/104 buses, two RS-232 serial ports, and sockets for 256 KB of battery-backed SRAM.

Most '386 signals are available via two 72-pin connectors at each end of the module. Processor address space is divided into eight regions, with four chip-select lines available at the connector for memory expansion or



accessing peripheral hardware.

Two interrupt controllers provide a total of 16 interrupt lines. Four are available at the connector, so additional interrupt controllers can be cascaded for real-time control applications.

One timer channel is internally connected to provide a standard timer interrupt. The other two are available at the connector for timing and

control signals. The built-in watchdog timer can reset the system, generate an interrupt, or indicate an alarm condition.

The connectors at each end of the card allow modules to be stacked together into a highly compact processing unit. Each module is 0.6" high with a stack pitch of 0.4". The complete system, containing processor, memory, power supply, two serial communication ports, network, and PCMCIA interfaces can be packed into a space just 1.8" high.

A base-configuration 25-MHz **i386** populated with 1-MB RAM and OS-MB flash ROM sells for \$265 in small quantities. The networking and PCMCIA interfaces sell for \$210 each.

Skylake Talix

Unit 3 Lulworth Business Ctr. • Nutwood Way, Totton
Southampton SO4 3WW UK

t44 (0)1703 666403 • Fax: t44 (0)1703 663730

#504

LCD TERMINAL

The LCD Terminal provides a low-cost user interface for machine automation. It emulates a standard ANSI terminal and supports software handshaking. A 16 x 2 LCD display with optional backlighting combines with a 16-keypad port to provide a total user-interface solution.

In addition, an expansion header provides the designer with an additional eight lines of 8-bit analog inputs or DIO. The additional I/O can be remotely accessed via ESC-command sequences.

The LCD Terminal ships with an RS-232 or RS-485 serial interface. The RS-485 version can be networked with up to 16 terminals per

network. The serial connection is made via a DE-9 male connector, and other connections are made via 0.1" center positive-locking connectors. Overall dimensions are 3.5" x 1.5" x 1.25".

The LCD Terminal sells for \$59 in 100-piece quantities. The single-quantity price is \$89.

Oklahoma Technical
Manufacturing Co.
2005 Exchange Ave.
Oklahoma City, OK 73108
(405) 524-2499
Fax: (405) 232-1833

#505

NEW PRODUCT NEWS

MOTION CONTROLLER

The DMC-1411 is a PC/104 form-factor, high-performance motion-controller card for 1 1/2 axes. It uses a 32-bit specialized microprocessor and controls both servo and stepper motors. Its 16-bit motor-command DAC and sophisticated PID filter gives precise control. A 250-line x 80-character program memory enables the control of custom applications without a host PC. Multitasking threads permit the simultaneous execution of two independent application programs. External sensors interface with seven uncommitted inputs and three uncommitted outputs to synchronize the controller with external events.

The DMC-1411 controls any motion, including jogging, point-to-point positioning, electronic gearing and CAM, and contouring. Its error-handling capabilities include hardware and


programmable software limits, automatic error shut-off, amplifier enable, user-defined error subroutines, and a watchdog timer.

An intuitive, easy-to-use, programming language that includes 126 user-defined variables plus arithmetic operations is provided. Multiple arrays enable real-time data capture of up to 1000 elements. A Visual Basic Tool Kit, which provides both Visual Basic Extensions (VBXs) and OLE Custom Controls (OCXs), is available to autotune, view trajectories on a PC monitor, and create powerful, application-specific operator interfaces with Microsoft Visual Basic.

The DMC-1411 motion-controller card sells for \$595.

Galil Motion Control, Inc.
203 Ravendale Dr.
Mountain View, CA 94043
(415) 967-1700
Fax: (415) 967-1751

#506



SBC WITH EMBEDDED SYSTEM CONTROLLER

The EPC-30 provides complete PC compatibility on a single board with standard ISA-bus and PC/104 connectors. It's based on the 25-MHz, 3.3-V Intel '386EX C-step processor for low-power applications. Optionally, the unit can be configured for 33-MHz operation at 5 V. It includes the PhoenixPICO BIOS, Phoenix PicoCard PCMCIA software, and PhoenixPICO Flash Array RFA software.


The EPC-30 features a RadiSys R380EX embedded system controller that provides a simple interface to additional chips (e.g., video or PCMCIA controllers). The R380EX also includes a DRAM controller, keyboard and mouse controller, real-time clock, enhanced IDE interface, and power-management features.

The EPC-30 has a wide range of optional features and is customizable for specific applications. Standard configuration includes 256-KB boot-block flash, and IDE- and floppy-drive interfaces. It also has a PC/104 expansion header, three RS-232 ports, parallel port, keyboard and mouse connectors, and watchdog timer. Optional features include up to 4-MB DRAM, plus a standard SIMM socket for an additional 1-16 MB. Other available options include a flash filesystem with 1, 2, or 4 MB of flash memory, LCD/SVGA video support for CRT and flat-panel displays, single-slot PCMCIA Type I/II/III support, and RS-422 and TTL ports.

In the standard configuration, the EPC30 sells for \$416 in quantities of 100.

RadiSys Corp.
5445 NE Dawson Creek Dr.
Hillsboro, OR 97124
(503) 615-1100
Fax: (503) 615-1150

#507



NEW PRODUCT NEWS


EMBEDDED MICROCONTROLLER

The Am186/88ER devices are highly integrated 'x86-compatible microcontrollers with onboard RAM. Shrinking the standard '186/'188 core through 0.35- μ m process technology allows cost-effective, on-chip integration of 32 KB of RAM, eliminating the need for external RAM. The microcontrollers deliver '386-level performance when coupled with inexpensive 70-ns flash memory, and the devices operate on only 3.3 V of power with 5-V tolerant I/O. They are ideally suited for a wide variety of memory-intensive applications in networking, telecommunications, mass storage, industrial control, and hand-held devices.

The devices feature a choice of 100-pin TQFP or PQFP packaging in 20-, 25-, 33-, and 40-MHz frequencies. Pricing for 10,000-piece quantities of the 20-MHz Am186ER in PQFP packaging is \$11.65. The 20-MHz Am188ER in PQFP packaging sells for \$11.05.

AMD
One AMD Pl.
Sunnyvale, CA 94088
(408) 732-2400 • Fax: (512) 602-7639

#508



REMOTE POWER CARD!

3 VERSIONS:

RESET

WATCHDOG RESETS PC IF IT HANGS FOR HARDWARE OR SOFTWARE REASONS

PHONE

TURN ON PC WITH PHONE, SHARE VOICE / MODEM LINE, CONTROL AC APPLIANCES


TIMER

WAKEUP OR SHUTDOWN PC, LATE NITE BACKUP / MODEM, CONTROL AC APPLIANCES

95\$

27\$ OEM

ALL MVS CARDS INCLUDE ASM, BASIC, AND C SOURCE FOR PC OR SBC



8 CHAN ADC

DATA ACQUISITION, SERVO CTL, AUDIO
8-BIT RESOLUTION 22KHZ SAMPLE RATE
SHARP CUTOFF ANTI-ALIAS FILTER
CREATE STEREO BLASTER (.VOC) FILES

95\$

2 CHAN DAC

VOICE MAIL, MUSIC, ALARMS, CTL VOLT
8-BIT RESOLUTION 44KHZ SAMPLE RATE
PLAYS MONO / STEREO BLASTER FILES
FUNCTIONS AS DIGITAL ATTENUATOR TOO

75\$



MVS BOX 850
MERRIMACK, NH
(508) 712 8507

5 YEAR LIMITED WARRANTY
PING IN USA

#104


Real-Time Embedded Systems

RTKernel

Professional, high-performance real-time multitasking system for DOS and 16-bit Embedded Systems.

For Borland C/C+++, Microsoft C/C++, and Borland Pascal.

Libraries: \$550 Source Code: add \$500



RTTarget-32

Cross Development System for 32-bit Embedded Systems.

Supports Intel 386 and higher, as little as 66K RAM/ROM.

For Borland C/C+++, Microsoft C/C+++, and Watcom C/C++.

Libraries: \$1700 Source Code: add \$1000




RTKernel-32

Professional, high-performance real-time multitasking system for 32-bit Embedded Systems.

Supports Intel 386 and higher.

For Borland C/C+++, Microsoft C/C+++, and Watcom C/C++.

Libraries: \$1950 Source Code: add \$1650



In North America, please contact:

On Time
88 Christian Avenue, Setauket, NY 11733, USA
Phone (516) 689-6654
Telefax (516) 689-1172
email info@on-time.com



Other Countries:

On Time Embedded Systems
Hohweg 49, 22083 Hamburg, GERMANY
Phone +49-40-2279405
Fax +49-40-2279263
email 102212.3101@compuserve.com

NO ROYALTIES! FREE DEMO DISK!
<http://www.on-time.com>

On Time

REAL-TIME AND SYSTEM SOFTWARE

FEATURES

14 Low-Power, Solid-State DC Motor Control

20 A Radio-Controlled Car and a 68HC11 Microprocessor

26 Using Microchip's PIC14000 Battery-Management Chip

34 Pulse'05 Pulse Generator

42 The Global Positioning System

Low-Power, Solid-State DC Motor Control

Finding mechanical switches passé? Well then, why not try National Semiconductor's LM1951 and LM2941? They do the job and provide low-power DC control.

FEATURE ARTICLE

Bill Payne



I've always liked robotics and industrial controls. It seems so magical to move things

in a precise manner.

In prior designs, I used mechanical relays for the power-control circuits. These relays controlled two different types of motors—some propelled a chassis while others steered.

However, mechanical relays produce a lot of heat and require a lot of current. I've tried semiconductor devices before and ended up with burn holes in my circuit board.

The traditional 2N3055 is a hearty device. But, the control circuits and protection circuits are sometimes more complicated than the devices they control.

With the advent of power MOSFETS, I thought my troubles were over. But, the original devices had about the same number of problems as the transistors they replaced.

While I was looking through a databook, I stumbled on the National Semiconductor LM1951 Solid-State 1-A Switch and the LM2941 1-A Low Drop-

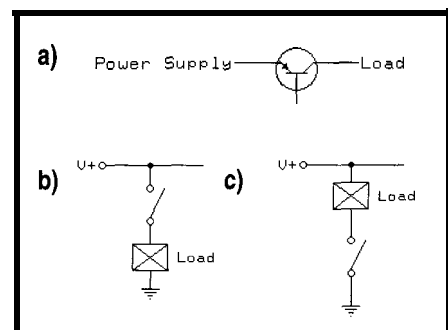


Figure 1—This figure depicts traditional PNP switch (a), load isolation with a high side switch (b), and the danger of energizing the load with a low side switch (c). The low side switch always has power applied to the load and needs only to find a ground to energize it.

out Adjustable Regulator. These devices meet the core requirements for a low-power DC-control system.

I can construct a system where one power source is used for the sensitive semiconductor devices and another for noisy motor environments. Power requirements are minimized through the use of optocoupled control circuits.

Power consumption in the quiescent state is at microamp levels since optically controlled voltage regulators allow multiple power supplies in the system.

In this article, I discuss this system with an optoisolated four-channel switch. Each switch is capable of handling up to 1 A. Current consumption is reduced by switching off the regulated power supplies to both sides of the four-channel solid-state switches.

LM1951 1 -A SOLID-STATE SWITCH

The LM1951 is a high-side switch designed for high-current, high-voltage applications. High-side switches control power to a ground-referenced load.

They operate similar to the older PNP transistor switch shown in Figure 1a. Since they are based on MOSFET technology, they don't have the inherent problems of the older transistors.

RDSon is much lower in these devices than in an older transistor switch. Therefore, the device is much more efficient and produces much less heat.

The high-side switch protects the load from being energized if the load wire is inadvertently shorted to ground (see Figure 1b). In contrast, the low-side switch applies power to the load and becomes energized if the load wire is shorted to ground (see Figure 1c).

The LM1951 can be continually connected to a live power source (e.g., a robot's motor-drive battery). This setup provides a high-side switch and




Photo 1-The LM2941 is heatsinked to prevent thermal shutdown. The blue connector at the bottom left side of the board is the drive for the optoisolator and controls the regulator's operation.

typically draws 0.1 μ A of current in the quiescent state.

An internal current sink connects to LM1951's input control pin (5), which pulls low if left open, thereby switching the device off.

The switch also has a built-in error-detection circuit. This circuit drives an output error flag when there's a fault condition on the load side (see Table 1). These detected error conditions can be things such as a stalled motor on a robot or an output short to ground.

The LM1951 drives both resistive and inductive loads (e.g., relays and indicator lamps). It has a built-in negative voltage clamp for the kickback from switching inductive loads. This clamp can handle approximately -30 V.

There is also built-in protection for reverse voltage on the load side of the device. Not that this would be a problem—we all use red wires for the positive side of the supply and black wires for the negative side—right?

The LM1951 also withstands transients up to ± 85 V on the load side, which is great for those applications

that have a lot of electrical noise on the supply load.

DC drive motors can wreak more havoc on semiconductor devices than you can imagine. The surge currents at startup and the fluctuations under varying loads work wonders on the supply rails to semiconductor devices.

Most importantly, the device can be driven by virtually any type of TTL or CMOS device. The input has built-in hysteresis in the 1.2-1.4-V range. Basically, it turns on at ~ 1.4 V and is guaranteed to turn off at 0.8 V and below.

The switching time for both turn-on and -off is typically 2 μ s. The V_{CC} supply to the part can be anywhere from 4.5 to 26 V, providing a wide range of applications for all types of resistive and inductive drives (e.g., lamps and motors).

A LOW DROPOUT REGULATOR

The LM2941 is a general-purpose positive-voltage regulator. It can source up to 1 A of current with a typical dropout voltage of 0.5 V and a maximum of 1 V over the entire operating temperature range. It has an internal quiescent-current-reduction circuit which reduces the device's ground-pin current.

This reduction circuit comes into effect when the differential between the input and output voltages exceeds -3 V. The quiescent current while driving

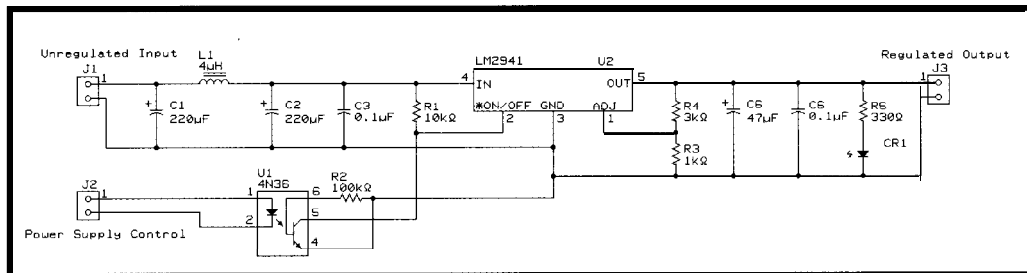


Figure 2—The control to the regulator is pulled up by the unregulated input voltage source. This guarantees that the regulator is in the off condition whenever the device is connected to a power source.

a 1-A load with an input-output differential of 5 V is only 30 mA. When the device is in the dropout mode ($V_{in} - V_{out} \leq 3$ V), higher quiescent currents exist.

The LM2941 was designed for vehicular applications. So, it has built-in protection for things like input line transients and reversal of battery input leads.

During a load dump on the input side of the regulator, the input voltage can momentarily exceed the specified maximum operating voltage of the regulator, so the regulator shuts down to protect the internal circuits and the load. It also has customary features such as short-circuit and thermal-overload protection.

The LM2941 can also be switched on and off from a TTL or CMOS device. It turns on and off at a threshold




Photo 2—Each of the four MOSFET switches at the top of the board is heatsinked to prevent thermal shutdown. The four outputs on the right side can drive devices such as OPTO-22s or any other load up to 20 mA each.

voltage of typically 1.3 V, meeting standard TTL requirements.

The maximum input supply voltage is 26 V. The output voltage is adjustable from 5 to 20 V.

POWER-SUPPLY MODULE

The power-supply module shown in photo 1 and Figure 2 is designed for electrically noisy environments.

The input circuit consists of an LC filter network made up of two 220- μ F capacitors and a 4- μ H choke. The two large capacitors and the choke smooth and filter the input voltage applied to J1. This provides a fairly clean input for the LM2941.

Output voltages are set by the resistor divider network R3 and R4. R3 is fixed at 1 k Ω to ensure that the input bias current error of the regulator's adjust pin is negligible.

R4's value is selected to provide a reference voltage of approximately 1.275 V. R4 is solved as:

$$R4 = R3 \left(\frac{V_{out}}{V_{ref}} - 1 \right)$$

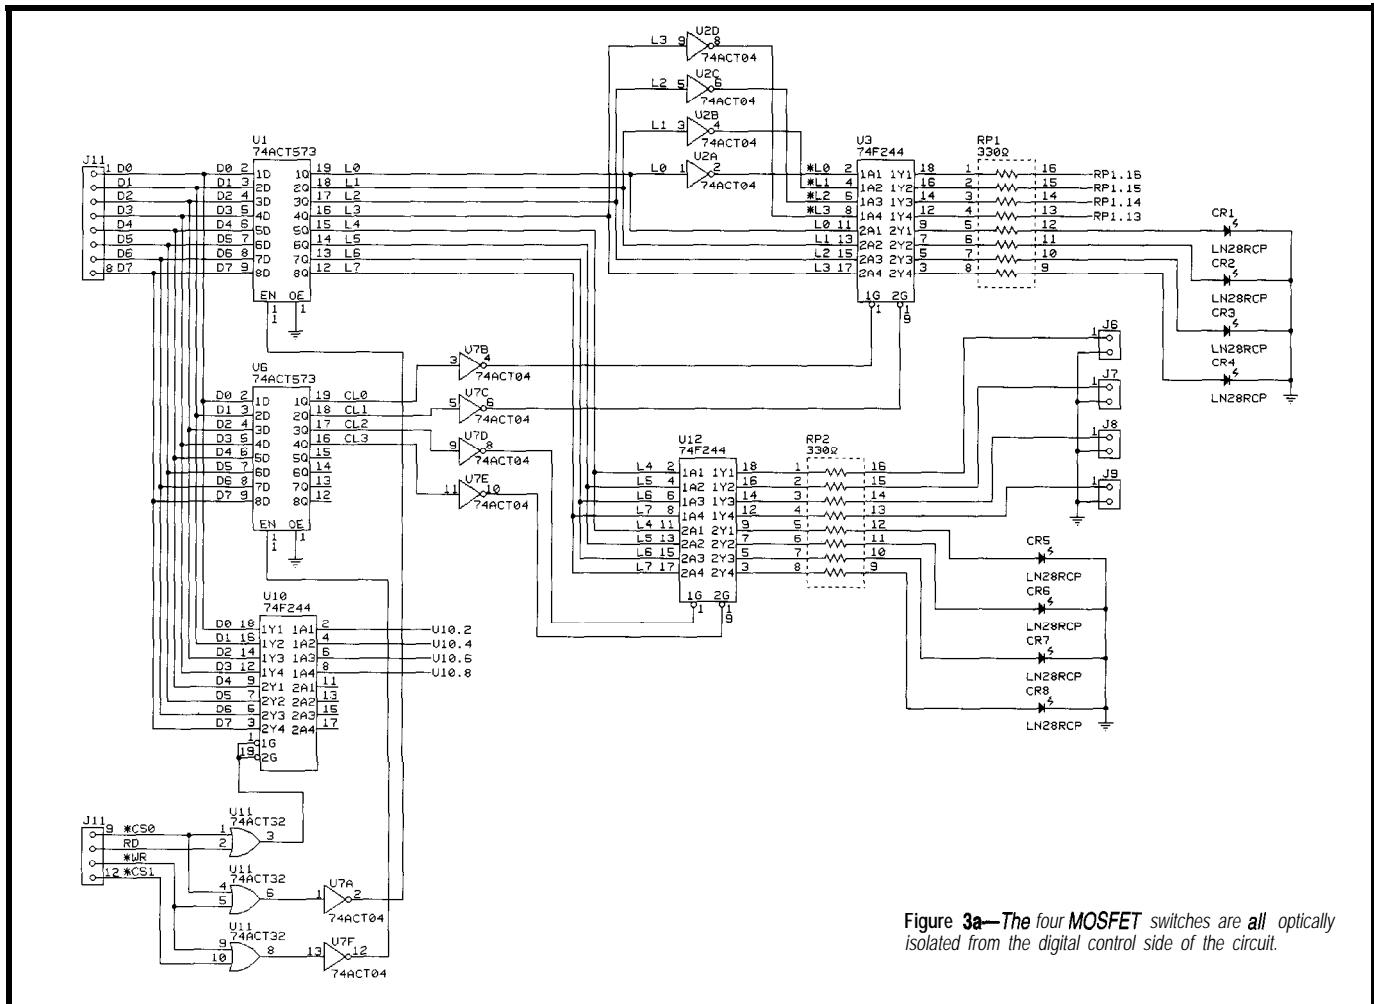


Figure 3a—The four MOSFET switches are all optically isolated from the digital control side of the circuit.

Resistor R4 is set at 3 kΩ in this design to provide a stable 5-V output on connector J3. Output capacitor C5 is set at 47 μF. It must be at least 22 μF to maintain stability.

A 0.1-μF capacitor is paralleled to C5 to catch any high-frequency transients which may have gotten through the regulator. J3 controls the LM2941.

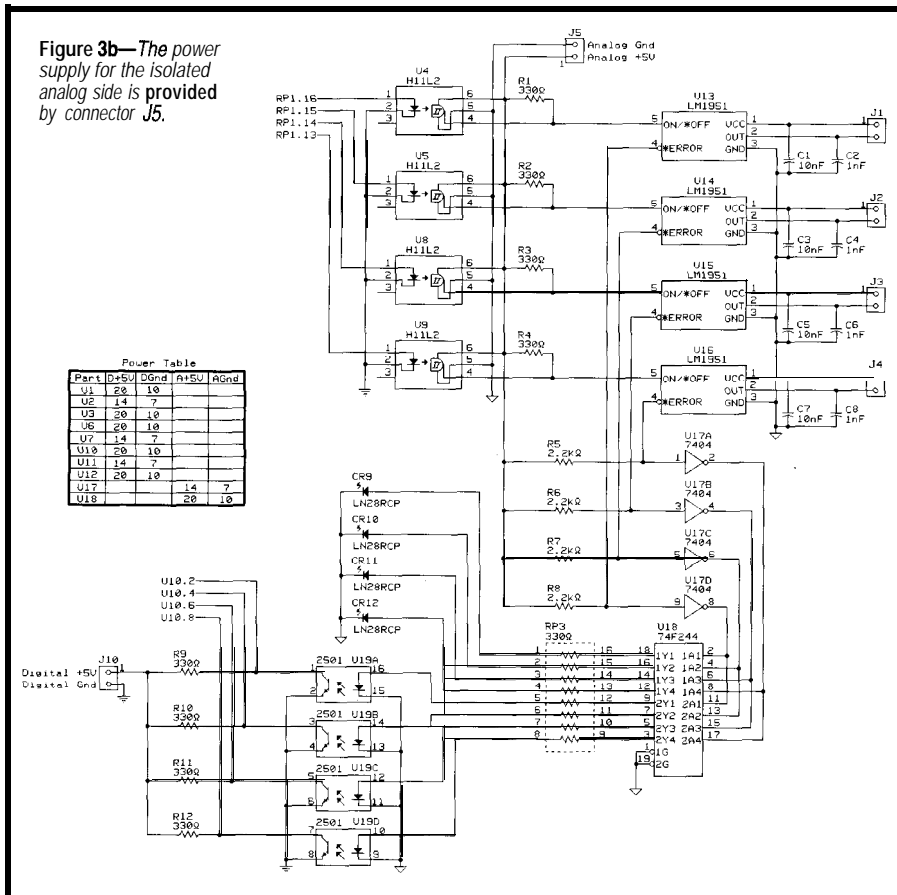
A standard 4N36 optoisolator provides isolation between the control and power circuits. Pull-up resistor R1 connects to the regulator's input side.

It pulls the regulator control pin (2) high when power is applied to connector J1 of the power-supply board. This guarantees that the regulator is turned off at powerup and that no output voltage is present on connector J3. J2 lets the regulator be switched on or off under control of a microprocessor.

The power-supply module provides power to separate boards in a system. Each board can be powered up after the processor performs initialization.

This reduces current consumption dramatically. Boards can be shut down individually to reduce drain.



Figure 3b—The power supply for the isolated analog side is provided by connector J5.



HAL - 4

EEG Biofeedback Brainwave Analyzer

The HAL-4 kit is a complete battery-operated 4-channel electroencephalograph (EEG) which measures a mere 6" x 7". HAL is sensitive enough to even distinguish different conscious states-between concentrated mental activity and pleasant daydreaming. HAL gathers all relevant alpha, beta, and theta brainwave signals within the range of 4-20 Hz and presents it in a serial digitized format that can be easily recorded or analyzed. HAL's operation is straightforward. It samples four channels of analog brainwave data 64 times per second and transmits this digitized data serially to a PC at 4800 bps. There, using a Fast Fourier Transform to determine frequency, amplitude, and phase components, the results are graphically displayed in real time for each side of the brain.



HAL-4 KIT NEW PACKAGE PRICE - \$279 +SHIPPING
Contains HAL-4 PCB and all circuit components, source code on PC diskette, serial connection cable, and four extra sets of disposable electrodes.

to order the HAL-4 Kit or to receive a catalog,
CALL: (860) 875-2751 OR FAX: (860) 871-041 1

CIRCUIT CELLAR KITS • 4 PARK STREET
SUITE 12 • VERNON • CT 06066

* The Circuit Cellar Hemispheric Activation Level detector is presented as an engineering example of the design techniques used in acquiring brainwave signals. This Hemispheric Activation Level detector is not a medically approved device, no medical claims are made for this device, and it should not be used for medical diagnostic purposes. Furthermore, safe use requires HAL be battery operated only!

4-CHANNEL SOLID-STATE SWITCH

The four-channel solid-state switch module in Photo 2 provides four 1-A optoisolated high-side switch outputs to any DC load under 26 V. It also provides for very low power consumption while in the Off state.

Numerous LEDs indicate status throughout the design. The schematic is shown in Figures 3a and b.

Connector J11 acts as an interface from the module to a microprocessor-based system. This generic interface consists of data lines (D0–D7), chip selects (*CS0 and *CS1), a read line (*RD), and a write line (*WR).

By asserting ● CS0 and ● RD, the data at the inputs to the 74ACT244, U10, is available on the J1 1-bus interface. These inputs are optocoupled to the error-flag outputs of each LM1951. They indicate the various fault conditions defined in Table 1.

Each error flag connects to an LED on the isolated high-voltage side of the module as a visual status indicator. A fault condition turns on the associated LED, and it becomes readable by a processor connected to the J1 1 bus.

By asserting ● CS0 and ● WR, the data on the J11 bus is written to the 74ACT573, U1, which is a transparent latch. The lower four nibbles of the latch data (D0–D3) control the LM1951.

Writing a logic 1 to these bits turns on the associated switch. In addition, the 74F244, U3, replicates the data on

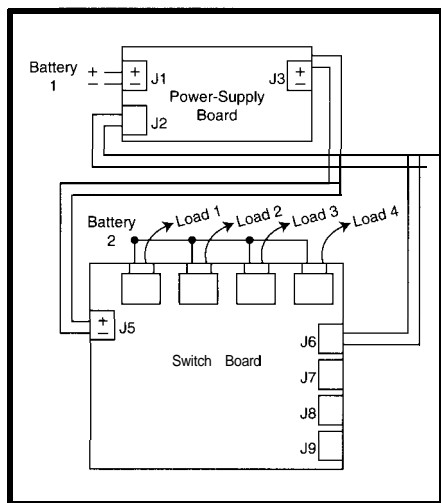


Figure 4—The power-supply module can connect to the 4-channel MOSFET switch module for ultra-low power consumption. The analog section of the 4-channel MOSFET switch is powered by a controllable power supply, as is the digital side of circuit, allowing the system to turn off when not needed.

Fault Condition	$V_{on/off}$	V_{out}	Error Flag
Normal	L	L	H
	H	H	H
Overvoltage	L	L	L
	H	L	L
Thermal Shutdown	L	L	L
	H	L	L
V_{out} Short to Gnd	L	L	H
	H	L	L
V_{out} Short to V_{CC}	L	H	L
	H	H	L
Open Load	L	L	H
	H	H	L
Current Limit	L	L	H
	H	H	L

Table 1—Various fault conditions can be sensed by turning the LM1951 on and off and reading the error flag. Conditions such as a stalled motor, open circuit, and supply short to ground can be determined by a microprocessor and visually displayed.

DO–D3 and drives an associated LED that indicates the state of each switch.

The upper nibble of the latch data goes to another 74F244, U12, which drives four current-limited outputs at connectors J6–J9. It also drives four LEDs used as status indicators for the output state on each connector.

By asserting *CS1 and WR, the data on the J1 1 bus is written to the other 74ACT573, U6, which is a transparent latch. D4–D7 are not used in this design. DO–D3 control the outputs of the two 74F244 buffers at U3 and U12.

By setting bit D0 to a logic 1, the drivers to the optoisolators are enabled, so data in the transparent latch at U1 can pass to the optoisolators. In this way, the data can be set up before the switches are activated. This idea is good in a system that contains drive motors (e.g., a robot).

Bit D1 enables or disables the LED indicators for solid-state switches. If you need to conserve power, turn the LEDs off. The output of the latch is inverted, so a logic 1 enables the buffer.

By setting bit D2 to a logic 1, the drivers for connectors J6, J7, J8, and J9 are enabled. The data from the upper nibble of the transparent latch at U1 then flows to the outside world.

Bit D3 is used in the same manner as bit D1 is for the buffer at U3. It enables or disables the visual status-indicator LEDs for the four driven lines at connectors J6, J7, J8, and J9.

Connector J10 is the digital supply input for the module. It powers all of the bus-interface circuits and the opto-driver circuits. It must be present before power is supplied to other parts of the module.

Connector J5 is the digital supply for the high-voltage switch side of the module. It should be powered up after the latches are stable and contain the appropriate data.

The loads connected to the LM1951s remain off due to the internal pull-down on the control pin. This ensures no power flows until it's appropriate.

Each connector at J1, J2, J3, and J4 is connected with pin 1 to the power source for the solid-state switches. Pin 2 of each connector is tied to the load as is standard for a high-side switch.

CONTROLLED CURRENT

The four-channel solid-state switch module combined with multiple power-supply modules allows for a robust control system that features very low quiescent power consumption. Figure 4 depicts one of many configurations for this system.

This combination of modules lets a microprocessor control current consumption. For example, if the drive motors in a robot do not need to be energized, associated control circuits can be turned off to conserve the battery. Current consumption in this state is in the 1- μ A range for each switch. □

Bill Payne has many years' experience as a digital design engineer. He holds two semiconductor patents and has others pending. You may reach Bill at bpayne@ionet.net.

SOURCES

LM1951, LM2941
National Semiconductor
P.O. Box 58090
Santa Clara, CA 95052-8090
(408) 721-5000

I R S

401 Very Useful
402 Moderately Useful
403 Not Useful

FEATURE ARTICLE

Dean Hoyt

A Radio-Controlled Car and a 68HC11 Microprocessor

Automated devices are becoming critically important for safety as well as convenience. Dean shows how he added a 68HC11 micro to a radio-controlled model so it could record and then retrace its path.

Question—what do you get when you mix radio-controlled (RC) model cars, a microcontroller, and some software? Answer—a project that expands the abilities of the model car.

My aim was to enable the model to record a path traversed and then retrace the recorded path. To keep the model free to run around, I also wanted to control the path recording and playback from the transmitter.

This sort of radio-controlled device has applications beyond mere entertainment, however. Autonomous robots are “trained” to go into hazardous sites on clean-up or retrieval missions.

But, for my project, I first had to understand the car’s control system.

FIGURING IT OUT

As you see in Photo 1, my RC car is a 1:12 scale model of a Lamborghini Cheetah, which is basically a four-wheel-drive Jeep. Mine, however, doesn’t have four-wheel drive.

The Cheetah’s control system is a radio receiver with two servos and a battery pack. One servo controls the steering of the front wheels, and the other controls the speed of the electric motor.

The power train is an electric motor with a separate battery pack. The motor attaches to a gear box and differen-

tial to drive the wheels. I modify the top speed and torque by changing a gear in the gear box, which means disassembling the gear box. Because the Cheetah isn’t used outside, I left it in the fast-speed, low-torque gear.

The electric motor’s speed is controlled by voltage applied to the motor. I altered this voltage by having the servo motor move a wiper on a variable power resistor—the slower the speed, the more resistance added to the circuit.

It’s not the best way to control a motor’s speed, but it’s inexpensive. Fancier motor-speed controllers take in the signal that goes to the speed-control servo and controls a bank of power transistors for speed and motor direction.

To enable a microcontroller to monitor the servos’ positions for recording a path, a couple of options come to mind. You can measure the position of the control arm attached to the servo motor, or you can monitor the control signal going to the servo motor.

Since I wanted to add as little hardware to the system as possible and I planned to control the servos when playing back a recorded path, I decided to monitor the control signal to the servos.

To determine just how the radio receiver controlled the servo motors, I hooked up a scope to the three wires going to the servo and discovered a very simple control system.

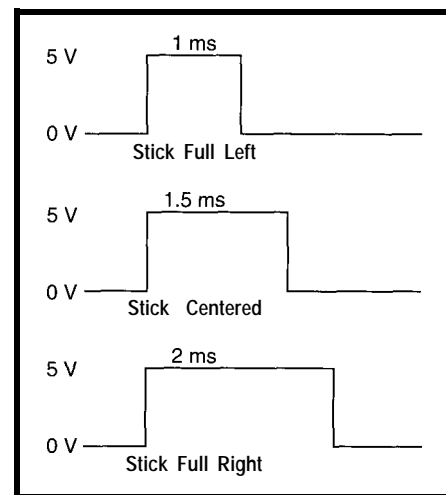


Figure 1—The servos in model cars and airplanes are controlled by a PWM signal. Here are the pulses for various stick positions. Any signal source generating the proper waveform can control the servos. I use the radio receiver and the 68HC11 to produce the waveforms.




Figure 2—This composite pulse is the output of the analog electronics in the receiver. This waveform is sent into a chip that parcels out pulses to each channel. Using this signal simplifies the hardware interface to the radio.

Two wires carry power to the servo (of course), and the third wire carries a PWM signal to the servo. The servo is made for digital control from a micro because the PWM signal is 0-5 V.

My radio has six channels. The pulse going to the servo has a width of 1-2 ms depending on the position of the radio receiver's control stick. With the control stick dead center (its default position), it produces a 1.5-ms pulse, as shown in Figure 1.

The transmitter has two joysticks on it [each controlling two servos] and two switches (each controlling one servo). The switches move the servos from one extreme to the other. The receiver controls up to six servos. The pulses for each servo are offset in time.

Inside the radio receiver, I found some analog parts (the RF receiver and signal conditioning) and one IC. The output of the IC turns out to be the signal to the individual servos.

Through probing, I determined the IC's input to be a multiplexed version of all the outputs to the servos pre-

ceded by a 3-ms sync pulse. The length of one information packet (for my 6-channel receiver) was 20 ms. Your receiver may vary.

The IC chip detects the sync pulse and then deals the pulses out to the receiver's outputs. The first pulse after the sync pulse goes to the first servo, second pulse to second servo, and so on (see Figure 2).

PLANNING

To add the new control to the existing vehicle, I selected Motorola's 68HC 11 processor. Since the peripherals are built into the processor, no additional hardware was needed to monitor the PWM signals controlling the servos.

My Motorola 68HC 11 evaluation board was set up with two serial ports, some additional RAM, headers for access to all of the 68HC11 pins, and a monitor program. Using an assembler on a PC, I generated an S19 object file and downloaded it to the EVB using Buffalo, the built-in monitor.

The 68HC 11's peripheral that appealed to me most for this project was its ability to measure and produce pulses. But, to understand this, you need some background information.

The 68HC11 has a free-running 16-bit counter, which is always on while the processor is powered up.

Eight special pins on the processor use this counter—three are inputs, and the other five are outputs. The three inputs can capture the value of the free-running counter in a set of 16-bit registers (one register per input pin) when a rising or falling edge occurs on the corresponding input pin.

The state of the output pins can be affected by a second set of 16-bit registers (one register per output pin). A 16-bit value can be stored in the output registers.

When this value equals the 16-bit free-running counter, the output for the matching register can be affected. The effect is programmable (i.e., set output pin high, set output pin low, or toggle state of output pin).

The inputs can produce interrupts when they capture the counter's value. The outputs produce an interrupt when the output register matches the free-running counter.

Each interrupt is vectored, so there's no need to determine its source. This special hardware in the 68HC11 makes measuring and producing PWM signals easy and accurate.

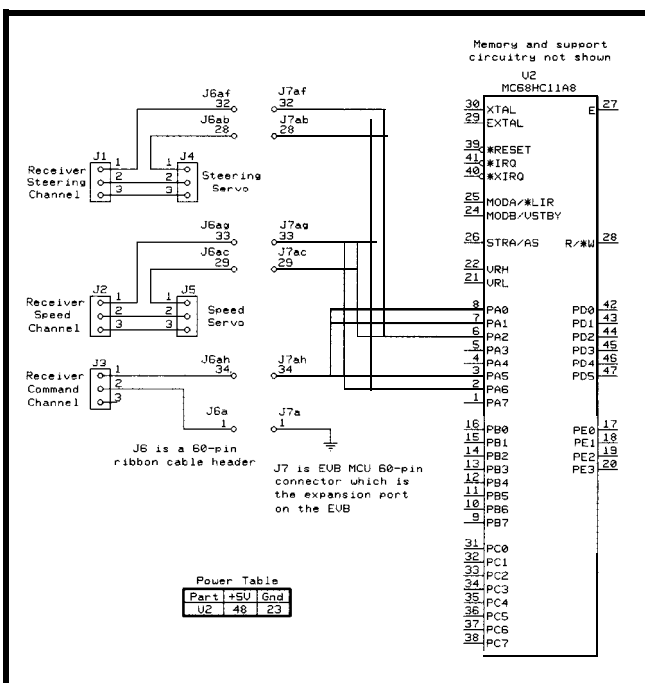


Figure 3—This schematic shows the connections of the radio and servos to the 68HC11. Otherwise, the project boils down to software and cables.

HARDWARE INSTALLATION

Once I had the hardware to measure and produce the control signals for the servos, it was time to figure out how to insert the hardware into the controls.

I thought about installing analog switches between the servos and radio receiver. With the switches closed, the radio would directly control the servos and the 68HC11 would record signals sent to the servos. To control the Cheeta, the switches would open (isolating the radio from the servos), and the 68HC 11 would produce control signals for the servos.

After pondering this option a while, I decided to leave the 68HC11 in the loop all the time. It can then monitor signals from the radio and produce the control signals for the servos.

This method has its advantages. External hardware, analog switches, and their control lines aren't needed. As well, there's no longer the need to sync up the changeover from radio control to processor control and back. If the 68HC11 is always in the loop, there's no sync problem.

This final installation is shown in Figure 3.

USER INTERFACE

To talk to the 68HC11 via the radio, I chose to monitor a third channel from the radio (one of the switches on the transmitter) and have two operating modes (command and control). With the toggle switch away from the user, the Cheetah is in command mode. If the switch is toward the operator, it's in control mode.

When the processor is powered on, it assumes command mode and waits for the transmitter's switch to reenter command mode. This allows a safe startup given the unknown state of the switch on the transmitter.

After this, the Cheetah follows the switch and changes modes whenever the switch is moved. One joystick on the transmitter controls steering, speed, and direction (forward and reverse).

In command mode, there are four commands-transparent move (stick right), learn (stick left), replay path forward (stick forward), and replay path backward (stick backward).

Once a command is issued, the car waits for the switch to change to control mode. In transparent mode, the 68HC11 is an expensive wire and the joystick controls the vehicle's direction and speed.

In learn mode, the joystick still controls the Cheetah, but the 68HC11 records every pulse that goes to the servos.

In replay, the joystick does nothing. After replay is selected, the Cheetah waits for control mode to be entered and then replays the previously recorded series of pulses to the servos, effectively replaying the path it followed.

One added software feature is that the recorded path isn't destroyed as it is played. The path can be replayed forever.

Listing 1--The pulse-width signal from the radio is received by the 68HC11 using its built-in port A Input Capture hardware. Two interrupt routines are used for each signal measured. The first one notes the time of the rising edge and sets up to capture the falling edge. The second routine notes the time of the falling edge, calculates how long the signal was high, and stores this. It then sets up to capture the rising edge.

```
* Interrupt routine to start timing pulse on ICI
SV2IC1r          Processes rising edge of pulse
LDD #SV2IC1f    Change to falling-edge service routine
STD PVIC1+1
LDD TIC1        Save time of leading edge
ANDB #MASK      Remove jitter
STD FRSTE1
LDX #REGBAS     Reconfigure ICI for falling edge
BCLR oTCTL2,X $30
BSET oTCTL2,X $20
BCLR oTFLG1,X $FB Reset the interrupt flag
RTI

* Interrupt routine to finish timing pulse on ICI
SV2IC1f          Process falling edge
LDD #SV2IC1r    Change to leading-edge service routine
STD PVIC1+1
LDD TIC1        Get time of interrupt
ANDB #MASK      Remove jitter
SUBD FRSTE1     Subtract time of leading edge to get
STD ICIHI       time it was high. Store this

TST LEARN       Should interrupt routine control servos?
BEQ SKIP1      No branch
STD OFFHI2     Yes, so set servo high time
LDD PWMPIP
SUBD ICIHI
STD OFFLO2     Low time = period - high time
SKIP1          Reconfigure ICI for rising edge
LDX #REGBAS
BCLR oTCTL2,X $30
BSET oTCTL2,X $10
BCLR oTFLG1,X $FB Reset interrupt flag
INC ICI1DONE   Inform main routine that new value captured
RTI
```

SOFTWARE INTERRUPTS

I wanted my software simple-just five interrupts. I had one for each radio channel (command switch, steering servo, and speed and direction servo) and one for each servo.

To simplify the interrupts, I wrote ten different routines-one for each rising and each falling edge.

In Listing 1, the program detects a rising edge on the input and calls the rising-edge interrupt routine. When the rising edge is detected, the 16-bit free-running counter's value is captured and stored. The input pin is reconfigured to capture a falling edge, and the interrupt routine is switched to the falling-edge routine.

I thought it would be quicker in code time to change the interrupt routine than figuring out which edge is currently being worked on and put the branch statements into the interrupt routine.

When the falling-edge routines are called again, they capture the 16-bit free-running counter and calculate the

number of counts that the pulse was high (i.e., falling-rising-edge time) and store this value. They also set a flag indicating that a new value was captured.

If the Cheetah is in either transparent or learn mode, the falling-edge routine also stores the new value into the memory used by the output pulse-production routines.

The output pulse-production routines work in basically the same manner as the input routines. An initial value for output compare is programmed in, and the interrupt routine is set to a rising-edge routine. When the 16-bit free-running counter matches the initial value, the pin goes high and the interrupt routine is called.

The interrupt routine calculates when the falling edge occurs, configures the hardware to produce a falling edge, and changes the interrupt routine to a falling-edge routine. This routine calculates how long the output is low by subtracting the high time from 20 ms. This calculation allows the

output waveform's period to match the input waveform's period.

Other transmitters may have different periods, and a future enhancement would be to measure the actual period of the input waveform. Once the falling-edge routine goes low, it sets the next match value for the 16-bit counter, configures the hardware to produce a rising edge, and calls the rising-edge routine.

I wrote one master set of pulse-width-measuring routines and one master set of output pulse-width-production routines. These routines handle one servo channel.

The second servo-channel interrupt routines differ only in the variables and hardware used. The two servo channels correspond to the steering and speed servos. The last set of input interrupt routines (for the command switch channel) still measures the input pulse width.

The difference is that the routine just determines which side of 1.5 ms the pulse is on and sets a command/control flag to represent the correct state. The command/control variable tracks the state of the transmitter's command/control switch.

I then enabled the program to decode commands and store the values captured and thus learn a path. Decoding was pretty easy.

When the time value for one axis on the joystick exceeded a threshold, a command was selected. The thresholds are set almost to the extremes the joystick can go to, which effectively eliminates the chance of receiving false commands.

The first threshold exceeded then picks a command. To select another command, command mode has to be exited and reentered. The main program monitors the flags that indicate when a new sample is captured.

With a new sample, the thresholds are checked. A command is selected, the proper routine jumped to, and the state of the system alters to reflect the command (i.e., recording, replaying, or transparent).

FITTING IT ALL IN

There's one other small detail. I have a limited amount of memory—




Photo 1—Here are all the hardware components for a trainable RC car. Just add the software.

about 8 KB of RAM—and I want to store two 16-bit pulse lengths every 20 ms. The best-case scenario gives me about 40 s of recording time—not very long to play with an RC car.

Although this recording time was much too short, it didn't make sense to install more memory on the EVB board. Instead, I compressed the path information to save memory.

Looking at the data produced by the input pulse-width capture routines, I determined that the four LSBs were more noise than data. With the four bits masked off, the data for unchanging servo positions wasn't altered, so I masked them off and implemented a run-length compression as shown in Listing 2.

For the compression, the number of times the same value occurs consecutively is noted and stored along with the value. Now, I'm using eight bytes (two two-byte counts and two two-byte times) for each time sample that's different along the path.

This means that when the control stick is not moving, the counter is incremented, but no more memory is consumed. The counter is not checked for overflow. With a 16-bit counter, it's assumed that nobody would hold the stick in the same place for more than 1300 s (i.e., 21 min.).

In normal operation, the stick isn't moved for straight runs and constant turns. With run-length compression implemented, recording time changes

Listing 2—To maximize the limited RAM in the 68HC11 system, use a simple run-length compression. If the current value to store equals the last value, increment a count. Otherwise, add the new value to the list and set its count to one. For slow-changing data like the signals to the servos, this method works well. The information stored is the high time of a pulse.

```
* New value was recorded by input compare routine for Channel 1
RECORD1 CLR IC1DONE
        LDD IC1HI          Load new value
        CPD 0,X           Compare with last pulse width
        BNE NEWREC1
        LDD 2,X           Add 1 to pulse count
        ADDD #1
        STD 2,X
        BRA WAITL2
NEWREC1 INX              Store new value in new location
        INX
        INX
        STX MAXREC1
        STD 0,X
        LDD #1           First value stored
        STD 2,X
        BRA WAITL2
```

from its former *best* case of 40 s to a *worst* case of 40 s-not too shabby for a simple compression method.

The absolute best case would be if the driver moved the stick once every 20 min. You'd have a recording time of 14 days but not such an interesting path. Since the drive-motor batteries last 2030 min., the compression gives more than enough path-recording time.

PLAYING

Now, the system's complete! I have an RC car with onboard smarts that can be controlled strictly from the radio transmitter, and it can record a decently long path.

The acid test was showing it off. I drove the Cheetah across the room, around one of my friends, and back to me while recording the path.

Then, I placed it in its starting position and told it to replay the path. Everybody was impressed when it drove across the room, around my friend, and back to me.

Some potential expansions include measuring the period of the signal so

any transmitter and receiver could be used or pulse-width modulating the motor with some FETs rather than using a variable resistor.

You could also find the composite signal in the radio receiver and feed this one signal into the 68HC11, rather than each servo's signal individually. The 68HC11 would then access all of the radio's channels, and you'd only need one input pin to the processor.

If you're interested in this sort of application, you can get started with one of the many 68HC 11 single-board computers advertised in *INK*. This project needs no other external hardware.

So, whether your primary goal is maintaining safety in hazardous situations, cleaning your house more efficiently, or simply having fun, you can effectively combine hardware, software, and radio-controlled devices in your designs. Your only limit-your imagination. ☒

Dean Hoyt is a computer engineer in Imation's Systems Lab. He currently

designs and implements embedded control solutions for research projects. You may reach Dean at dchoyt@imation.com.

SOFTWARE

Code for this article is available via the Circuit Cellar BBS, the Circuit Cellar Web site, or Software on Disk for this issue. Please see the end of *ConnecTime* for downloading and ordering information.

SOURCES


68HC11, 68HC11 EVB

Motorola
MCU Information Line
P.O. Box 13026
Austin, TX 7871 1-3026
(512) 328-2268
Fax: (512) 891-4465
<http://freeware.aus.sps.mot.com/>

I R S

404 Very Useful
405 Moderately Useful
406 Not Useful

Touch The Future



LCD Touch Monitors
LCD Touch Screens
VGA LCD Displays
LCD Controllers
ISA, PC 104, Analog, Video



EARTH

Computer Technologies

lowest Prices on Earth!

27101 Aliso Creek Rd - # 154 - Aliso Viejo - CA - 92656
Ph: 714-448-9368 - Fax: 714-448-9316

Email: oemsales@flat-panel.com

FREE CATALOG available at <http://www.flat-panel.com>


#108

Announcing: **FlashICE** NOW AVAILABLE!

FlashICE, the world's first Flash emulator, is the only tool available for developing and testing embedded Flash update routines.

- ❖ Emulates Flash Memory
- ❖ User Control of internal Flash Functions
- ❖ Verifies programming algorithm
- ❖ Profiles sector write or erase cycles

Web: www.gei.com



Grammar Engine Inc.

921 Eastwind Dr., Suite 122
Westerville, OH 43081
614/899-7878
Fax 614/899-7888


#109

FEATURE ARTICLE

Fred Eady

Using Microchip's PIC14000 Battery-Management Chip

Fred introduces us to Microchip's new mixed-signal PIC14000. It has all the features of a PIC but also 192 bytes of RAM, 4 KB of EPROM, and analog peripherals. It's a lot of complexity in a flexible package.



I'm a digital kind of guy. But, from time to time, I attempt to speak a little analog. I'll bet you can, too. Come on, say "PIC." Roll your tongue, "P-I-C. . ."

You're probably thinking, "This guy's lost his mind! Everybody and their sister knows what a PIC is. . .and it ain't analog! Hey, Fred, can you say 'op-amp'? You've got your signals crossed, buddy." [Raucous laughter.]

Well, since you asked, yes, I can. . . PIC...{rolling my tongue}. .P-I-C.. My signals aren't crossed. They're mixed.

In the old days when micro names were short and started with a 6, 8, and Z, melding analog and digital required fancy engineering and a bunch of parts. But, now it's the late twentieth century, and there's a silicon-bending corporation out there called Microchip that's changing the way we use and think about microcontrollers.

So, stop PIC'n on me, and let's mix it up with the new mixed-signal PIC-14000.

THE PIC14000

Based on the PIC16C74, this member of the PIC family has all trademarks of a PIC. Throw in a ton of RAM (192 bytes) and EPROM (4 KB) along with some analog peripherals, and there it is—the PIC 14000.

If you've read any of my stuff, you know I like to statistically describe databooks. Why change now! The PIC14000 databook is exactly 0.325" thick when properly compressed. Add 0.217" (single-sided) for the application notes, and you have a little over a half inch of paper chock full of newly created PIC analog and digital stuff.

For the truly mobile and tree-conscious, the 0.048" thick 1996 Microchip Technical Library CD provides

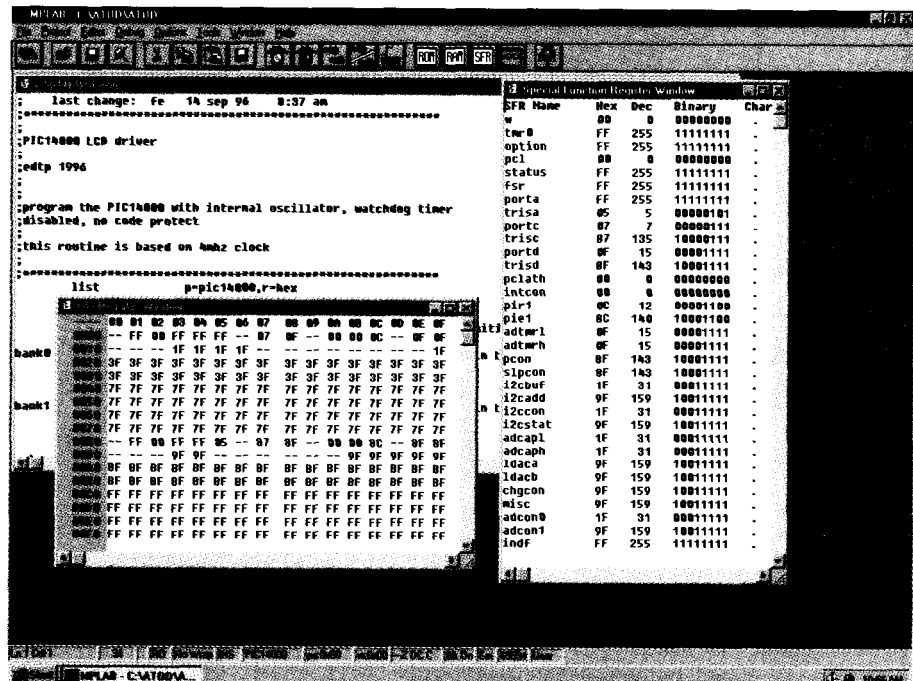


Photo 1—Microchip's PICMaster puts us in the cockpit of the PIC14000. Lots of knobs to turn here!

the same information (sans new app notes) via your PC or Mac. You can get the very latest scoop plus application software from the Microchip Web site.

Basically, the PIC 14000 consists of three major subsystems which include the CPU core along with the analog and digital peripherals. We're all familiar with PIC CPU operations, so let's concentrate on the features that make the PIC14000 stand out.

BITS TO VOLTS

Some real brainionic person once conjectured, "I think, therefore I am." Likewise, you're reading this, therefore you're literate.

So, I'll spare you *Fred's Datasheet Digest* of the PIC14000's digital-to-analog circuitry. You want fine hairs, you comb the datasheet.

Instead, let's take a look at how the PIC14000 DAC circuitry works by putting it to the test. Where I can, I'll use Microchip's Pro Mate, PICMaster, and MPLab as a window to the PIC-14000's internals.

I'll end up with a piece of PIC14000 hardware that measures voltages, controls FETs, and talks to me via an LCD and PC. Photo 1 is a programmer's eye view of the PIC14000.

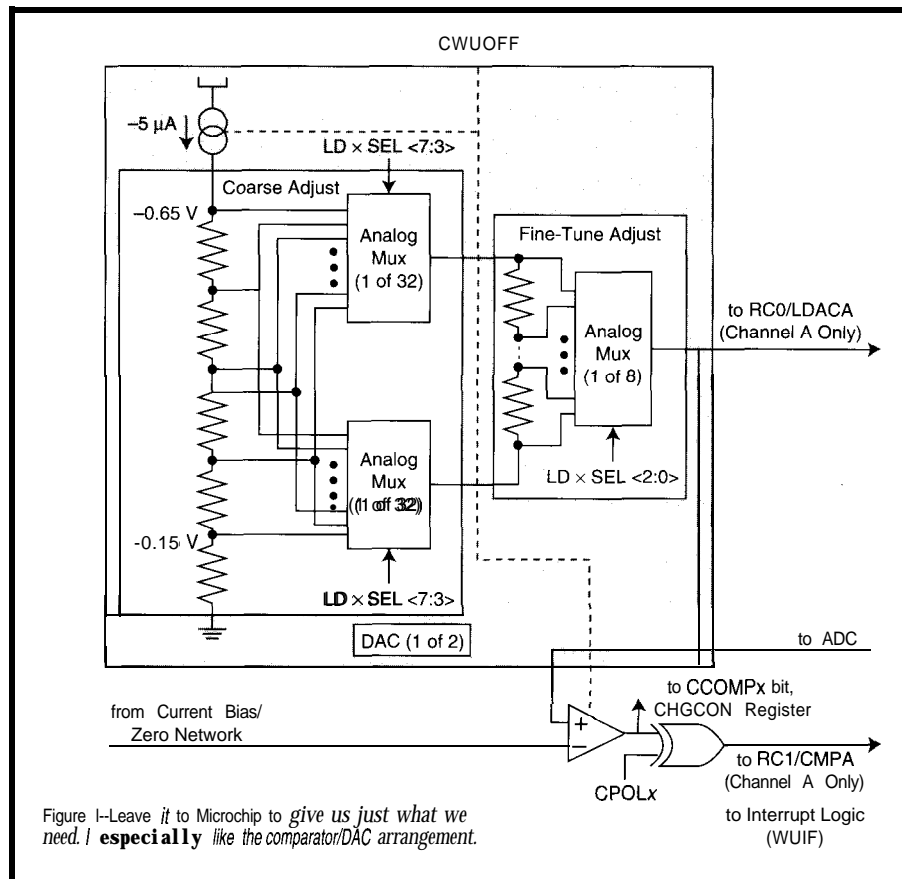
As the section header implies, the PIC14000 contains a DAC. Actually, there are two. They enable the PIC-14000 to generate analog voltages depending on the value loaded into the DAC's control register.

The DAC voltage range is 0.15–0.85 V. This synthesized voltage can drive an op-amp or set the trip point for the pair of comparators glued to the DAC circuitry. The DAC/comparator combination allows current sensing and control depending on how you logically connect the parts.

The DACs are called LDACA and LDACB. Each DAC is made up of a coarse resistor ladder controlled by a five-bit analog multiplexer. The ladder is driven by a constant-current source.

The coarse-ladder output is buffered and fed into a three-bit mux that fine-tunes the DAC output. Both muxes are addressable via registers.

To make things interesting, the output of each DAC is fed to the inverting input of a comparator. Yeah, a



comparator! The DAC output is also fed to the ADC, which enables the DAC to set the comparator trip point while being monitored by the ADC. The noninverting comparator input is connected to a bias/zero network that is pinned out.

This arrangement is a hotbed for all sorts of current-sensing and voltage-monitoring applications. You could even fabricate a couple "electronic pots" with the help of some good FETs.

The LDACs are structured in a way that enables the comparators to be tripped on charge or discharge currents. This current-detect mode can be configured to interrupt the PIC14000 or wake it up from sleep. It's also possible to configure the comparator inputs to form a window with the voltage envelope set by each individual LDAC.

The PIC14000 also performs charge control. In this mode, the comparator's output is tied directly to the gate of a FET. The DAC is configured to set the comparator trip point, and the noninverting comparator input senses current via a sense resistor.

This arrangement continually attempts to make the voltage across the

sense resistor equal to the trip voltage set by the DAC. Figure 1 gives a simplified view of the LDAC. Pretty slick, huh? My application uses the DACs to form the voltage window I spoke of.

VOLTS TO BITS

There exists the yin, and thus exists the yang. Unlike the PIC 16C74's successive approximation ADC module, the PIC14000's ADC function is done via slope conversion.

A capacitor is hung off the constant-current source CDAC pin. This perpetuates the voltage-versus-time capacitor theory and creates a voltage ramp. The A/D conversion is performed by measuring the time it takes the voltage across the cap to equal the input analog voltage being converted.

The PIC 14000 does this by incrementing a counter as the voltage ramps up. When the voltages are equal, the count at that instant is stored in 16-bit format for later use.

That's nice.. .but how do you know what the count equates to as a voltage? Easy. The PIC14000 is endowed with an internal voltage reference called the bandgap reference.

Thus, the input analog voltage is determined by two conversions. The first is done against the input analog voltage, and the second is performed on the internal bandgap reference. The input analog voltage is equal to the ADC count of the input voltage divided by the ADC count of the bandgap voltage multiplied by the absolute bandgap-reference voltage.

You're wondering where the heck did I get that "absolute" bandgap voltage value. No problem. It's stored in the PIC14000 calibration area.

Each PIC14000 is loaded with factory-set calibration constants located in a special area of the EPROM. There's also a calibration constant value for the A/D slope reference ratio, the temperature-sensor voltage and voltage slope, the internal oscillator, and the watchdog timer.

A bunch of physical things can mess up the linear ramp measurement-dielectric absorption and comparator offset voltage, for instance. So, I have to compensate for them if I want an accurate reading.

And how the heck do you do that, Fred? You guessed it-the PIC14000 has upper and lower slope reference voltages. Converting using these and mixing in the calibrated slope reference ratio gives an offset count.

This offset value is found by first subtracting the low-reference ADC count from the high-reference ADC count and multiplying the difference by the A/D slope reference ratio constant. The calculated result is subtracted from the low reference count to give the offset count.

This muddies our original voltage equation. So, let's just work it out. The voltage at any input pin is defined as:

$$V_{in} = \frac{N_{in} - N_{offset}}{N_{bg} - N_{offset}} \times K_{bg}$$

where N_{in} is the ADC count for the selected input, N_{bg} equals the ADC count for the bandgap reference, N_{offset} is the calculated offset count, and K_{bg} equals the absolute voltage value of the bandgap-reference voltage.

Sometimes, it's necessary to measure voltages below the usable range of the PIC14000 onboard ADC equipment. In fact, the PIC14000 can mea-

sure negative analog input voltages by providing a couple of special pins connected to a summing junction that adds 0.5 V to the analog input voltage before it gets to the ADC comparator.

Again, two conversions-one with the summing junction connected and the other without-are needed to produce a meaningful voltage reading. The difference of the two readings becomes the actual input voltage.

Before leaving the A/D module, let's talk about calibration constants. Analog and drift are constant companions.

For the most accurate measurements possible, each PIC14000 is factory calibrated. The resultant calibration

A/D slope reference ratio	K_{ref}
Bandgap reference voltage	K_{bg}
Temperature sensor voltage	V_{thrm}
Temperature sensor voltage slope	K_{tc}
Internal oscillator frequency multiplier	F_{osc}
Watchdog timer time-out	F_{wdt}

Table 1-These values are factory installed and unique for each PIC14000. Woe to the programmer with an itchy EPROM-eraser finger.

constants are stored in a reserved portion of its EPROM. The program uses this to minimize the effects of temperature, time, voltage, and process.

Be aware that you can erase these constants. Before you do any of your own PIC 14000 applications, carefully adhere to the warnings and suggestions in the app notes. They tell you how to preserve this precious calibration data.

Each value, with the exception of the internal oscillator frequency multiplier and watchdog-timer time-out values, is stored in 32-bit floating-point representation. Photo 2 is a Pro Mate view of the part I use in my application. Table 1 cross references the calibration values found at the bottom of Photo 2.

As you've seen, K_{ref} and K_{bg} are used in A/D conversions. V_{thrm} and K_{tc} are used with the internal temperature sensor.

F_{osc} is the single-byte value of the measured internal oscillator frequency. F_{osc} is normally used to compensate the internal clock for applications that require precise timing.

F_{wdt} is another single-byte value which is the watchdog time-out in milliseconds without using the post-scaler counter.

There's an entire app note devoted to the calibration constants (AN621) and a corresponding note for the floating-point representation (AN575). Microchip supplies a floating-point library for the PIC 14000 in AN575—thank goodness!-so I won't expound on that subject here.

Now, where are we? We know that the PIC14000 is a versatile mixed-signal processor. It can measure voltages, sense currents, produce voltages, and take your temperature.

We also know environmental variations encountered are compensated for via embedded calibration constants. That's nice.. .but what can it do!

APPLYING THE PIC14000

The PIC14000 datasheets and app notes tilt heavily toward battery-management applications. It's well-justified because the PIC14000 is power skimp and modular. This modularity enables you to power down portions of the PIC14000 not in use.

It also automatically powers down idle analog circuitry without user intervention! There's even a pin that lets the PIC 14000 (with the help of an external FET and resistor) regulate the incoming supply voltage from the battery it's monitoring for its own power.

Bottom line: it takes some digital trickery and the right mix of analog components to do battery stuff. The PIC14000 is perfect in this role.

I've always been infatuated with digital things pretending to be analog, and I have lots of old NiCd battery packs. So, here's what we're gonna do.

I CAN'T REMEMBER...

I hate NiCd battery "memory" arguments. If I have doubts about a battery pack, I discharge it to its minimum capacity and recharge it. Then, I use it. If it doesn't meet my expectations, I repeat the process. Eventually, I either gain or lose confidence in the pack and act accordingly.

Unfortunately, the discharge cycle is done with the piece of equipment the pack resides in. The no-no is that, usually, the piece of equipment with the suspect pack has battery management built in and the pack doesn't get fully cycled.

Let's use the PIC14000 and a handful of other stuff to perform a constant-current discharge on a NiCd battery pack down to a predetermined voltage level and shutoff. I'm doing this because I suspect the pack is not operating at its peak. Since it's a constant-current discharge, why not provide signaling so a PC can run timing and capacity calculations at the same time?

Since the PIC14000 is a PIC, I'll use its PIC power to monitor the pack voltages on an LCD and communicate with the PC program, too.

THE HARDWARE

As Figure 2 shows, the circuitry is very simple due to the highly integrated peripherals of the PIC14000. The LCD is attached in four-bit mode to conserve I/O pins, and two lines are allocated for bit banging to the PC parallel port.

Note there's no oscillator or crystal. The PIC14000 has an internal clock.

It seems the most complicated part of my design lies within the LM3900 circuitry. It's really a very simple ramp-and-hold configuration.

The first gate of the LM3900 provides a very low input bias current to the input of the second gate. The bias gate eliminates the need for FETs and the special-purpose integrated circuits normally needed for sensitive ramp-and-hold circuits.

The PIC14000's internal DACs and comparators sense the voltage across the current sense resistor, R5. If the current (voltage) reading is within the limits of the window, both comparators remain at TTL low which, in turn, puts the LM3900 circuitry in hold.

If the voltage violates the window in either direction, the associated comparator applies a TTL high to either the up or down side of the ramp-

and-hold module. The voltage at the output of the second gate ramps up or down accordingly.

This voltage drives the gate of the FET and makes it appear as a variable resistance. The rest is Ohm's law. By monitoring the preset voltage window and adjusting the resistance in the discharge path, I create a constant current across the FET/sense resistor combination. Is that clever or what?

The typical PIC14000 bandgap voltage hovers around 1.2 V. Scale input voltages to as close to the bandgap as possible to increase accuracy.

The voltage divider across the battery pack serves this purpose. The voltage the ADC module sees is then multiplied by the scaling factor before being displayed on the LCD.

My design allows for a discharge current of up to 300 mA, which is governed by the sense-resistor value and the DAC's voltage range. A smaller sense-resistor value results in a higher operating amperage range.

The PIC 14000 DACs and comparators are configured in charge control mode, so I can couple the DAC output voltages to the comparator inputs. I'm also using the bias/zero network, so the inputs are biased by 0.5 V.

There are as many ways to keep time as you can imagine. I could have included a real-time clock IC to measure elapsed time. But, it takes a lot of code to get the data into human form. If you're tight on EPROM space, forget it.

I'm not yearning for EPROM area in this project, and there's always an idle

PC lurking about. So, I usually reach over and put that little twerp to work.

Think about it. The video is already there.. the timer is already there.. the I/O is already there.. Get my point?

The timing idea here is simple. When the discharge cycle begins, the PIC14000 signals the PC via the parallel port to begin the elapsed time counter. When the battery pack reaches the cut-off voltage, the PIC 14000 signals the PC to stop the count and calculate the battery capacity.

The only input the PC needs is the discharge current. Using the elapsed time and discharge current, you can calculate the battery capacity in milli-ampere-hours.

THE SOFTWARE

Piece of cake. It's regular old PIC code. Although the PIC14000 has some fancy trappings, the code is the stuff you're used to:

1. Initialize the PIC14000 peripherals for the application
2. Read and store the summing voltage
3. Read and store the voltage across the sense resistor
4. Read and store the pack voltage
5. Read and store the bandgap voltage
6. Read and store the slope-reference voltages
7. Do steps 2-6 a minimum of 8 times and average the readings
8. Calculate the ADC offset voltage
9. Convert the voltage readings into BCD
10. Display the readings on the LCD

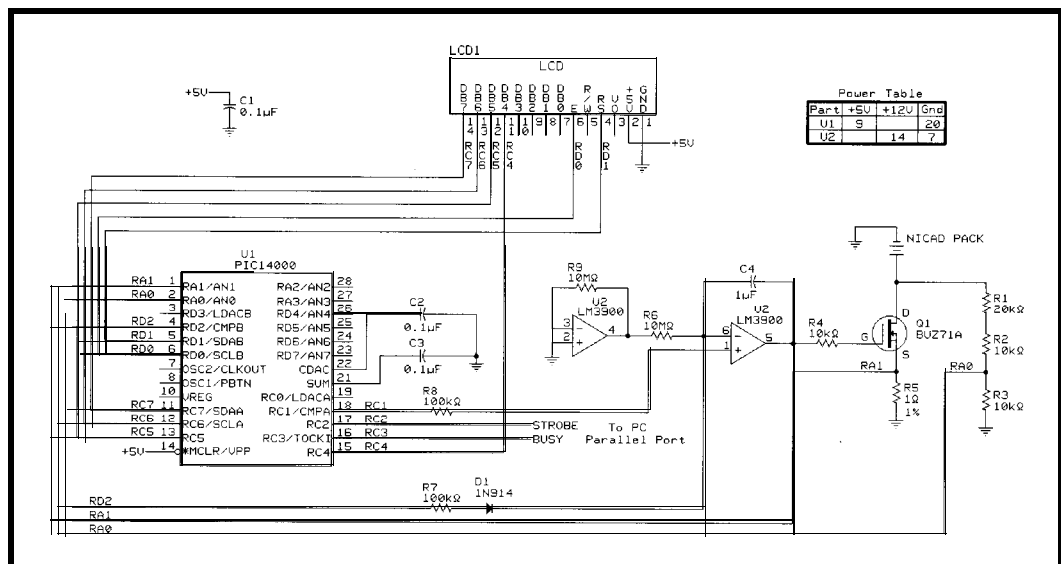


Figure 2—The PIC14000 can be used with an LCD display and a handful of parts to intelligently discharge a battery pack to a predetermined voltage level.




Photo 2—Once we've test-flown the PIC14000 application, Microchip's Pro Mate loads the cargo. Note the calibration constant window.

11. Check the pack for the desired cut-off voltage and act accordingly
12. Loop to 2 until the pack is at cutoff

If you're wondering where the code is to control the voltage window, don't hurt yourself. There is none.

In step 1, I load the LDAC registers to set the comparator's trip points. Once this is done, no CPU resource is necessary. The DACs operate like hardware DACs, and the comparators operate like discrete comparators. The PIC14000 CPU can concentrate on doing math and refreshing the LCD.

TAKE IT FURTHER

I could go on. The PIC 14000 also can communicate via I²C. It can even hibernate and awaken via interrupt.

It's a very complex device, but in reality, it's a flexible platform to launch applications from battery monitoring to environmental control.

Since the PIC 14000 can monitor negative voltages—perfect for 4–20-mA applications—it could be the heart of a loop-powered digital panel meter.

Thermocouple systems could save a sensor by using its internal temperature sensor. With ADC resolutions reaching 16 bits, it suits a multitude of sensor applications including pressure, humidity, position, and level sensing.

INK readers can produce working product using only a rock and a battery. Just think what you'd do with a PIC14000 and a battery! □

Thanks to Microchip's Hartono Darmawaskita for regulating my voltages and temperature during this article.

Fred Eady has over 19 years' experience as a systems engineer. He has worked with computers and communication systems large and small, simple and complex. His forte is embedded-systems design and communications. Fred may be reached at edtp@ddi.digital.net.

REFERENCES

- B. Dellacroce, "PIC 14000 A/D Theory and Implementation," Microchip Technology App Note AN624, DS00624B, 1996.
- R. Fisher, "PIC14000 Calibration Parameters," Microchip Technology App Note AN621, DS00621B, 1996.
- T.M. Frederiksen, "The LM3900: A New Current-Differencing Quad off Input Amplifiers," National Semiconductor App Note AN-72, 1991.
- Microchip Technology Inc., PIC14000 Datasheet, DS40122A, 1995.
- Microchip Technology Inc., PIC16/17 **Microcontroller Databook**, DS00158A, 1995.
- M. Rosenfield, "PIC14000 DAC Theory and Implementation," Microchip Technology App Note AN622, DS00622A, 1996.
- F. Testa, "IEEE 754 Compliant Floating-Point Routines," Microchip Technology App Note AN575, DS00575A, 1994.

SOURCES

PIC14000, PICMaster, Pro Mate, MPLab

Microchip Technology, Inc.
2355 W. Chandler Blvd.
Chandler, AZ 85224
(602) 786-7200
Fax: (602) 786-7277
<http://www.microchip.com/>


Timeline, Inc.
2539 W. 237th St., Bldg. F
Torrance, CA 90505
(310) 784-5488
Fax: (310) 784-7590

IRS

- 407 Very Useful
- 408 Moderately Useful
- 409 Not Useful


SDI EPROM EMULATORS

Powerful Tools,
Reasonably Priced.



E1 \$199.00 U.S.

Supports EPROMs to 128K x8
27C64 to 27C010 (1MEG)



E4 \$249.00 U.S.

Supports EPROMs to 512K x8
27C64 to 27C040 (4MEG)

- Powerful PC software tools
- Full screen & command line modes
- Supports all data formats
- Software configurable
- NEW 85nS access time standard, 35nS optional
- NEW 32 PIN PLCC adapters \$65.00
- NEW 3.3V version, operates at both 3.3V and 5V
- Memory backup feature
- High-speed downloading (LPTI-3) with error checking and correction
- Non-intrusive CMOS LP design
- Chain up to 8 units -any configuration
- Compact size in hard protective case
- 1 year warranty & free software upgrades
- Discounts on 2+ units

SDI SCANLON
DESIGN
INC.

5224 Blowers St. Halifax, NS Canada
TEL: 902 425 3938 FAX: 902 425 4098


Toll Free in North America
800 352.9770

www.isisnet.com/oceana/SDI
Email: 71303.1435@compuserve.com

Pulse'05 Pulse Generator

FEATURE ARTICLE

Brian Millier



I suspect that many of you-like me—bring your work home. The fast-paced world of computer electronics requires us to do this to keep abreast of developments in the field.

In my home shop, I have a computer, a used 50-MHz HP scope, a digital meter, and a good set of tools. But, what happens when I need all the expensive test equipment that's available at work but not at home?

While it may be true that our diagnostic abilities are heightened when we work without a full complement of test equipment, it's unlikely we'll purchase these special-purpose instruments with money from our own pockets. It seems to me that there should be an in-between.

In this article, I present a very low-cost pulse generator that, in many applications, replaces the expensive pulse generator in most electronics instrumentation labs.

The project centers around two ICs. The first is Motorola's 68HC705J1A—one of the world's least-expensive user-programmable microcontrollers (\$1.70). The second is Advanced Micro Devices' AM9513APC system-timing controller. At \$30, it's more expensive, but it's the Cadillac of timers with six discrete timing blocks and generous prescaling and signal-routing options.

To keep costs down and maximize convenience, the generator's timing parameters are entered using two rotary encoders and displayed via a two-line LCD module. Depending on the bounty of your junk box, you can duplicate this project for about \$50-100. I also provide some hints on how to use

low-cost rotary encoders in place of pots or keypads in your digital projects.

PULSE'05 FEATURES

The Pulse'05 consists of three independent pulse generators. While each section can be individually programmed, they all share a common timebase. The generation of pulse signals can therefore have specific time relationships with each other, which can be useful in certain situations. Figure 1 is a block diagram of the project.

Each section can be programmed in increments of microseconds, milliseconds, or seconds. While micro- and milliseconds range from 2 to 9999, seconds only count from 2 to 99. All timing functions are based on a 1-MHz quartz crystal timebase.

The first and most basic section—the clock section—provides a TTL square-wave output with a variable period. The only parameters for this section are the time units (s, ms, or μ s) and the width of the high portion of the square wave. Since the period of a square wave is twice that of the high portion of the waveform, the period range of this section is 4 μ s to 198 s.

The second section—output 1—is more versatile. It provides a TTL output pulse, which can be a repetitive or one-shot pulse event, triggered by an external source or a push button.

The three programmable parameters are the repetition rate (only applicable in repetitive mode), the delay time from trigger, and the pulse width. The latter two must share the same units but are otherwise independently programmable.

When used in repetitive mode, the repetitive trigger signal is available at the trigger connector as well as the actual pulse output on output 1. Without this signal, the Delay parameter has no relevance. When using the one-shot mode, the external trigger signal is applied to the same trigger connector.

The third section—output 2—shares the same programmable parameters as output 1 but operates only in repetitive mode. With the addition of a switch and connector, output 2 can also operate in the one-shot mode, like output 1.

While output 2 is completely independent of output 1 in terms of pro-

Bring a lot of work home where you don't have much test equipment? Brian shows how to make a low-cost pulse generator composed of two ICs—Motorola's 68HC705J1A and AMD's AM9513APC.

gramming, they share the same timebase and are phase coherent. Therefore, pulse trains with certain timing and phase relationships can be generated.

Output 2 adds the capability of variable output amplitude and offset. This ability is useful for analog circuits where something other than the nominal TTL levels is required.

The output amplitude ranges from 0 to 4 V, and the whole pulse train can be offset or shifted from 0 by ± 4 V. The total signal swing is limited to within ± 7 V of ground.

I didn't bring out the 1-MHz timebase signal to a discrete output connector, but this can easily be done for a clock signal that's higher than the clock output's 250-kHz upper limit. Just use the spare U3b to add an output buffer to the output signal.

Setting all 12 programmable parameters in a way that's both quick and intuitive to the user is challenging. I ruled out a keypad/LCD display, which is normally my favorite I/O scheme.

In test equipment, it's often useful to continuously vary the period or width of a pulse and watch its effect on the device being tested or designed. A keypad is too slow since the parameter's numeric value and an enter key need to be pressed for each incremental change.

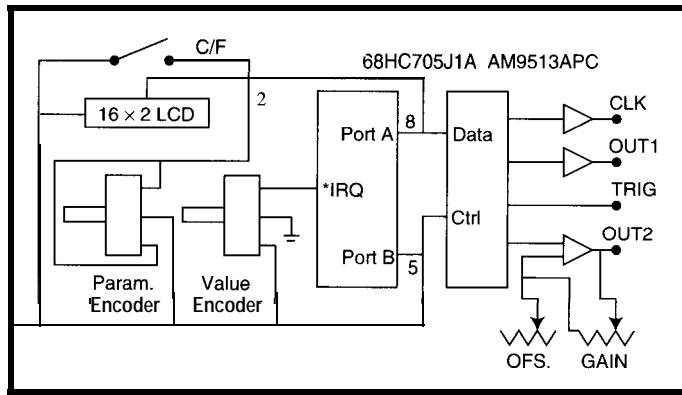


Figure 1—With only two major ICs, the Pulse'05 is a very simple design.

In an analog device, a potentiometer is ideal. But, in a digital device, it's too low in resolution and doesn't interface easily to a small microcontroller.

Bourns, a large manufacturer of potentiometers, also makes a low-cost rotary encoder (see sidebar "Using Low-Cost Rotary Encoders for Data Entry"). It resembles a potentiometer physically but provides two digital signals indicating how far the shaft has turned and in which direction.

This input device is ideal for this application, given the limited number of I/O pins on my microcontroller. Two such devices are used—one selects the parameter being varied, and the other changes the parameter's value.

For both fine adjustment of parameters and reasonably quick changes, an additional switch provides a coarse/fine toggle. In the fine setting, the parameter changes in one-unit steps. In coarse mode, the steps become 100

units. This momentary contact switch also serves as the entry device when a unit's parameter is selected.

TIMING IS EVERYTHING!

Since you've just read about the generator's features and ranges, you'll probably agree that neither the timer blocks present in common microcontrollers nor the ubiquitous and dirt-cheap 8253 PIT chip (found

in PC-compatible computers) can perform all of the necessary tasks. While they both include two or three 16-bit counters, they lack a truly comprehensive prescaling function and suffer from a small number of modes.

The best you can accomplish with either alternative is one fully programmable section. From a cost standpoint, it might make sense to use several 8253 programmable interval timer chips, but this involves much more wiring and requires four additional I/O lines that simply aren't available on the 68HC705J1A.

However, by choosing such an inexpensive microcontroller, I could spend the bulk of the project's budget on the AM9513APC. It's a superb timing chip for small-volume designs with complex timing functions. Although it's a 16-bit device, it's quite happy to work with an 8-bit bus as long as it receives the proper command on powerup.

Only ten interface lines are needed—which is very important since the 68HC705J1A has only 14 I/O lines.

Figure 2 shows a block diagram of the AM9513. The frequency-scaler block provides five outputs, all derived from the OSC frequency (F_{osc}).

F_{osc} is generated either by a quartz crystal connected between X1 and X2 or by a TTL-level signal fed into the X2 pin. I used a 1-MHz oscillator module fed into X2.

The divisor ratios available from the frequency scaler depend on the mode being used. Both binary and BCD modes are available, resulting in the divisor ratios in Table 1'.

The binary mode provides for a tremendous range of prescaling. However, due to limited program-memory

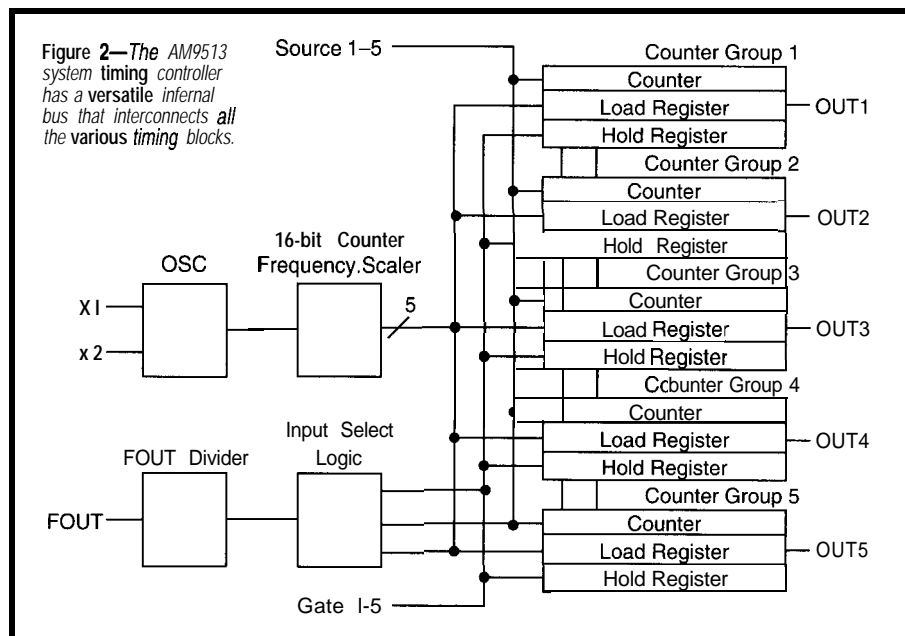


Figure 2—The AM9513 system timing controller has a versatile internal bus that interconnects all the various timing blocks.

Output	BCD Mode	Binary Mode
F1	OSC	OSC
F2	OSC/10	OSC/16
F3	OSC/100	OSC/256
F4	OSC/1000	OSC/4096
F5	OSC/10,000	OSC/65,536

Table 1--Both BCD and binary prescaler modes are available, resulting in the ratios shown here.

space in the 68HC705J1A, I chose the BCD mode for all 95 13 timer functions.

The timing functions are performed by the five blocks labeled counter group 1-counter group 5. Note that these blocks have lines dividing them into three sections to show that each block contains a 16-bit counter and 16-bit load and hold registers.

For all timing operations, the counter can be loaded from either the load or hold registers, or it can toggle between them. The latter option provides a way to generate PWM waveforms with each section of the waveform having a 16-bit resolution.

The following options add versatility to the timing blocks:

- the clock source for a block can come from the prescaler, the output of the next lowest numbered block, one of

five gate pins, or one of five source pins

- each block can operate in a recycle mode or perform only one timing operation per input trigger
- triggering can be either level or edge sensitive
- clocking can occur on either the positive or negative transition of the selected clock
- each block output can be programmed active high or active low or feed an internal flip-flop for a toggled output
- gating pins are available to gate the clock source

Having all these options provides 19 different timing modes (set. Table 2). While many are similar, there are enough to handle just about any timing chore you can dream up.

Although I seldom use this feature, all the blocks can be concatenated to form extremely long counting chains. For example, a real-time clock can be implemented with AM95 13 hardware. The chip also contains alarm registers that can be used in connection with such a scheme.

However, this is an expensive way to get an RTC. Give me a Dallas RTC

chip any day! I'll go into more detail about programming this chip when I discuss the firmware.

WHAT'S A MICRO TO DO?

Given the versatility of the STC chip, about the only thing left for the micro is to initialize the STC. This initialization is done by setting each section to the correct mode and loading appropriate parameters into the registers to reflect user-entered values.

This process may appear trivial, but it turns out that, like any user interface, a fair bit of code is involved. Indeed, I needed most of the 1240 bytes of EPROM in the micro.

Figure 3 is the schematic of the Pulse'05. Only four ICs are needed, thanks mainly to the STC's versatility and some micro interfacing tricks.

First, let's focus on the 68HC701J1A microcontroller (U1). It uses its own 4-MHz crystal (X1), which is internally divided by two to give an instruction cycle time of 0.5 μ s.

The micro needs to operate at its full 4-MHz speed to ensure adequate response to the encoders. For programming reasons, it made sense to use a 1- μ s timebase for the AM9513, so a

Using Low-Cost Rotary Encoders for Data Entry

Rotary encoders like the Bourns ECW1J-B24-BC0024, which costs about \$5, provide a viable alternative to potentiometers or keypads for data entry to small microcontrollers with limited I/O line availability. These units are roughly the same form factor as a potentiometer, use a standard 1/4" shaft, and can be PCB or panel mounted with a 3/8" nut.

This model is detented and has a resolution of 16 pulses per turn. The unit has two switches connected to a common terminal, resulting in a three-terminal interface, much like a pot. By grounding the common terminal and connecting pull-up resistors to each switch output, the output signals look like those in Figure 1 for each direction of rotation.

Assume the quadrature output 1 is connected to a negative-edge-sensitive interrupt on the micro and that quadrature output 2 is connected to any other available input port line. At the time represented by the dashed line, quadrature output 1 drops to zero, causing an interrupt to occur. Immediately on entering the interrupt service routine (ISR), the micro should check the state of the I/O port line connected to quadrature output 2.

If the state of the port line is high, the rotary encoder has been moved clockwise. The ISR need only set the

appropriate flags to indicate the presence and direction of that movement.

Since the switches in the encoder are mechanical, some contact bounce is present. At the falling edge of quadrature output 1, quadrature output 2 is always stable because it's halfway between transitions and no debounce is needed. If a 10-ms delay loop is placed in the ISR after it checks the state of the quadrature-output-2 port line, the debounce of the output-1 line is taken care of.

It's important to note that the micro's interrupt pin must be edge sensitive only, not level sensitive. If it's level sensitive, the ISR is reentered many times while the quadrature output 1 signal is low. This results in wildly incorrect readings and ties up the micro.




Figure 1--By grounding the common terminal and connecting pull-up resistors to each switch output, these signals are observed for each direction of rotation.

Table 2—The AM9513 STC has so many modes, there are barely enough letters in the alphabet to cover them.

Mode	Description
A	Software-triggered strobe with no hardware gating
B	Software-triggered strobe with level gating
C	Hardware-triggered strobe
D	Rate generator with no hardware gating
E	Rate generator with level gating
F	Nonretriggerable one-shot
G	Software-triggered delayed pulse one-shot
H	Software-triggered delayed pulse one-shot with hardware gating
I	Hardware-triggered delayed pulse strobe
J	Variable duty cycle rate generator with no hardware gating
K	Variable duty cycle rate generator with level gating
L	Hardware-triggered delayed pulse one-shot
N	Software-triggered strobe with level gating and hardware retriggering
O	Software-triggered strobe with edge gating and hardware retriggering
Q	Rate generator with synchronization
R	Retriggerable one-shot
S	No assigned description
V	Frequency-shift keying
X	Hardware save

common clock couldn't be used.

To interface to an LCD display, STC, encoders, and switches, you need some form of data bus. I used the '705J1A's Port A as an 8-bit data bus since I only needed the output direction.

The STC appears to the outside world as two registers. They are addressed using the C/*D pin, which is driven by the '705J1A's PBO line. The PB4 line rounds out the STC's interface line requirements by driving the *WR line. In this design, no reads of the STC registers are necessary, so the *RD line is tied high.

Also note that the high-order 8 bits of the STC's data bus must be tied high for the device to function properly in 8-bit mode! Luckily, the LCD

display has only two internal registers, which are also addressed via the same PBO line supporting that function for the STC. All that's needed to complete the interface to the LCD is an enable line, which is provided by PB1

The two Bourns ECW1J-B24-BC0024 encoders I used are shown to the far left of Figure 3. The value encoder is the one that is turned more frequently

and quickly, so it's connected to the *IRQ line and PB2. Using the *IRQ line, the micro hastily responds to any changes in signal level caused by shaft rotation.

By this time, I've almost run out of port pins on the '705J1A, and there's one rotary encoder and one switch left to connect. The response time needed in sensing the state of these two devices needn't be too fast since the operator must watch the LCD display and activate them accordingly to arrive at the desired state.

Therefore, I connected the two quadrature switch pins of the parameter encoder to DBO and DB1 and the coarse/fine toggle switch to DB2. The micro periodically goes through a mini keyscan routine by raising each of the three data bits in turn and sensing the switch's open/closed state using PB5.

Each line uses an isolation diode to prevent the data-bus lines from shorting when two or more switches close simultaneously. Up to eight switches can be sensed in this fashion using the whole eight bits of the data bus, but that wasn't necessary in this design.

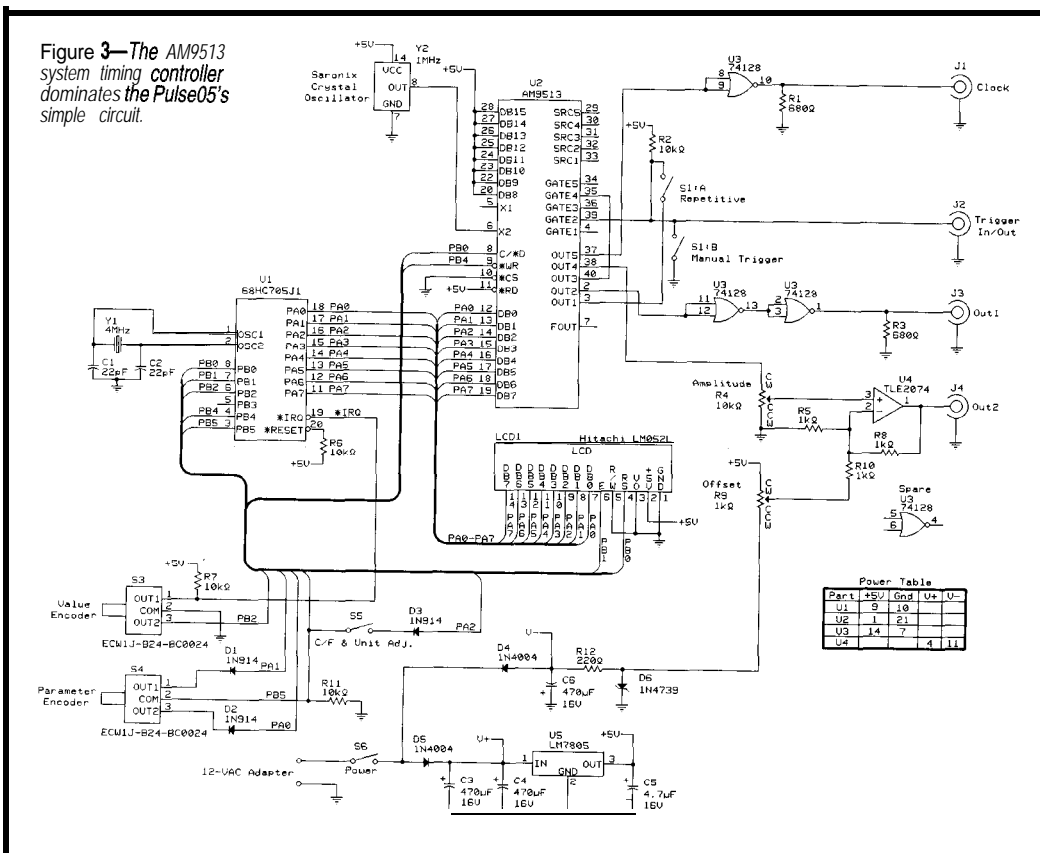
You may notice that PB3 is unused. At some point in an earlier project, I

must have fried this port pin on the one windowed (erasable) version of the chip that came with the Motorola 68HC705J1A evaluation kit.

You can't depend too heavily on the in-circuit emulation capability of the evaluator kit because it executes code hundreds of times more slowly than the actual micro using a '486 PC as the host. Therefore, I rely on the "burn micro, test circuit, erase micro" process.

So, I can't use this one port pin. But, I'm not inconvenienced enough to fork over \$40 for a new windowed chip-yet!

Each of the three pulse sections has a unique output circuit and, in the case of output 2, a trigger



Counter	Register		
	Mode	Load	Hold
1	FF01	FF09	FF11
2	FF02	FF0A	FF12
3	FF03	FF0B	FF13
4	FF04	FF0C	FF14
5	FF05	FF0D	FF15
Register			
Master Mode	FF17		
Alarm 1	FF07		
Alarm 2	FF0F		
Status	FF1F		

Table 3—These commands load the data pointer. The FF value in the MSB is hardwired into the circuitry due to the fact that DB-7-15 are all connected to V_{cc}.

option circuit as well. To buffer both the clock and output 1, I used a 74128 50-Ω driver IC.

I chose this rather than a more common TTL buffer IC to provide additional drive. Since the 74128 is a NOR gate and inverts the sense of the signal, two sections were used for output 1. The clock output is a square wave and, since inversion does not affect it, only one section was needed.

Output 2 connects directly to the amplitude pot, which feeds a TLE2074 high-slew-rate op-amp. This stage acts as both a buffer and a place to inject a variable voltage from the offset pot. The output signal is the low-impedance signal I described earlier.

Triggering for output 2 is more versatile than for the other sections. For the repetitive mode, STC OUT1 is connected as a trigger source for STC section 2 via the gate 2 input.

STC OUT1 provides a recurrent pulse at a rate decided by the user-entered period parameter. In one-shot mode, this recurrent signal is removed by opening switch S1a.

Instead, feed an external trigger signal through the trigger I/O connector, J2. Alternatively, a manual trigger can be applied by S1b, which is a momentary-contact section of switch S 1.

To keep the IC count down, no debouncing was done on this switch. This limits its usefulness to triggering pulses that are longer than the debounce time of the switch (generally in the 10-100-ms range, depending on the switch).

To further keep component count down, this switch is not isolated from the external triggering source con-

nected to J2. So, S1b should not be activated when an external source of triggering is connected.

The remaining circuitry in Figure 3 consists of simple plus and minus power supplies for the circuit. The unit operates from an external 12-VAC wall adapter, so the whole circuit fits into a small (8" x 6" x 1.5") enclosure.

THE FIRMWARE

The key to the AM9513 STC's versatility lies in its programmability. Interfacing the STC to a small micro is ideal because it requires only two control lines in addition to the 8-bit data bus. This requirement stems from the fact that the STC appears to the micro as only two registers—data and control.

The control register accepts commands and also sets the data-pointer register. Data is entered into the STC's 19 internal registers individually by first loading the data-pointer register with the appropriate value and then loading the data register with the actual data.

MM15: Scaler Control	0 = Binary, 1 = BCD
MM14: Data Pointer Control	0 = Enable Increment 1 = Disable Increment
MM13: Data Bus Width	0 = 8 bit, 1 = 16 bit
MM1 2: Fout Gate	0 = Fout on, 1 = Fout off (grounded)
MM1 1-MM8: Fout Divider	0000 DIV by 16 1000 DIV by 8 0001 DIV by 1 1001 DIV by 9 0010 DIV by 2 1010 DIV by 10 0011 DIV by 3 1011 DIV by 11 0100 DIV by 4 1100 DIV by 12 0101 DIV by 5 1101 DIV by 13 0110 DIV by 6 1110 DIV by 14 0111 DIV by 7 1111 DIV by 15
MM7-MM4: Fout Source	0000 F1 1000 GATE 3 0001 SRC 1 1001 GATE 4 0010 SRC2 1010 GATE5 0011 SRC3 1011 F1 0100 SRC4 1100 F2 0101 SRC 5 1101 F3 0110 GATE1 1110 F4 0111 GATE2 1111 F5
MM3: Compare 2 Enable	0 = Disable, 1 = Enable
MM2: Compare 1 Enable	0 = Disable, 1 = Enable
MM1-MM0: Time-of-Day Mode	00 TOD Disabled 01 TOD enabled: div. By 5 input 10 TOD enabled: div. By 6 input 11 TOD enabled: div. By 10 input

Table 4—To set the Master Mode Register, you must generate a 16-bit word whose bits reflect the options you desire.

Table 3 shows the command value that must be sent to the control register to set the data pointer for each of the 19 internal registers. For rapid register programming, AMD provides several sequencing modes that auto-increment the data pointer after each word is loaded, eliminating the need for constant reloading. I don't use this mode, instead choosing to prepare parameters in my own order and set the data pointer accordingly.

For the complete STC setup, see the section "95 13 STC Loading Routines" on the BBS. I'll describe only the setup of the Master Mode register and timer sections that generate the output-1 waveform.

Table 4 shows the Master Mode bit assignments. Although some features of the STC are unused in this design, the Master Mode register produces:

Scaler Control = BCD

Data-Pointer Control = Disable Increment

Data-Bus Width = 8

F_{out} Gate = On

F_{out} Divider = DIV by 10

F_{out} Source = F5

To achieve this, the data pointer is set to Master Mode register by sending a \$17 to the control register. Then, two bytes are sent to the data register—a \$F0 followed by a \$CA.

The STC requires data to be entered LSB first and then MSB. Therefore, the Master Mode register is loaded with \$CA F0, which meets the above criteria.

The Pulse'05 output- 1 signal is derived from STC counter groups 1 and 2. In repetitive mode, counter group 1 sets the repetition rate. In one-shot mode, this is performed by an external signal or the trigger push button.

Depending on the timing units chosen, the clock input to counter 1 is either F1(μs), F4 (ms), or F5 (s). Counter 1 is programmed for Mode D, which is set to count down in BCD mode repetitively with no gating, using the rising edge of the clock.

Also, it must produce a positive pulse at TC (terminal count). This positive pulse at the user-selected rate triggers counter 2, which produces a pulse with a user-set delay and width.




Photo 1—This view of the front panel of the Pulse'05 shows the LCD and parameter rotary encoders.

TIME'S UP!

Operating the Pulse'05 is straightforward. At powerup, after a sign-on message, the LCD displays the screen for setting the Clock High Period, as shown in Photo 1. Screens for setting all other parameters resemble this one.

On the second line to the right is a C or F, which is the coarse/fine adjust indication. C/F/U nits toggles between adjust increments of either 1 or 100 units. Adjust the value encoder to obtain the desired value.

Using the parameter encoder, you can run through all 12 screens. For every parameter, there are both value and units screens (other than delay and width, which share the same units).

When a units screen is displayed, only C/F/U nits has any affect. It toggles through the s/ms/ μ s selections.

The only inconsistency in this user interface is in the clock section. In all other cases, period defines the repetition time. For the clock section, the high portion of the square-wave output is specified because this output uses the TC-toggled mode of the STC.

Don't forget to set the Delay parameters for outputs 1 and 2. Each timer section must be loaded with a value greater than 1. Even if you don't need a delay, enter a 2- μ s value or the pulse will not be correct.

Good luck to anyone who decides to expand their home/shop with a Pulse'05. And, if you plan to use the AM9513 for the first time, save yourself some time by carefully reading the Technical Manual.

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. You may reach Brian at brian.millier@dal.ca.

SOFTWARE

The complete source code for this article is available via the Circuit Cellar BBS, the Circuit Cellar Web site, or Software on Disk for this issue. Please see the end of ConneCTime for downloading and ordering information.

SOURCES

68HC705 J1A microcontroller, Development Kit, and Technical Data manual, MC68HC705J1A/D

Motorola
P.O. Box 13026
Austin, TX 7871 1-3026
(512) 328-2268
Fax: (512) 891-4465
<http://freeware.aus.sps.mot.com/>

AM9513APC Timing Controller

Jameco
1355 Shoreway Rd.
Belmont, CA 94002-4100
(415) 592-8097
Fax: (415) 592-2503

AM9513A Technical manual 034021)

Advanced Micro Devices One AMD Pl.

Sunnyvale, CA 94088
(408) 749-5703
AMDlit@gomez.amd.com

Programmed 68HC705J1A

Brian Millier
3 1 Three Brooks Dr.
Hubley, NS
Canada B3Z1A3

ECW1J-B24-BC0024 encoders

Bourns, Inc.
1200 Columbia Ave.
Riverside, CA 92507-2114
(909) 781-5500
Fax: (909) 781-5273


IRS

410 Very Useful
411 Moderately Useful
412 Not Useful

EMBEDDED DOS CONTROLLERS AT 8051 PRICES

Use Your PC Development Tools

No MORE CRASH & BURN EPROM
Technology – 512 k FLASH



\$195
Qty 1 Price

Flashlite


DOS Single Board Computer
with 572 k FLASH Memory disk drive

- ✓ 10 Mhz/8 Mhz CPU
 - ✓ 512 k bytes RAM
 - ✓ 512 k/256 k FLASH
 - ✓ 2 Serial Ports
 - ✓ 24 Parallel I/O Lines
 - ✓ 2 Timers
 - ✓ 4 Interrupt Lines
 - ✓ 8 Analog Inputs
 - ✓ X-Modem File Transfer
- INCLUDES DOS & Utilities

USE YOUR TURBO C COMPILER OR
QUICKBASIC COMPILER
SAVE TIME, MONEY AND HEADACHES

A/D Converter

\$95
Qty 1 Price




- ✓ 8 Channels, 12 Bits
- ✓ 6 μ s. Conversion Time
- ✓ Clock/Calendar Option
- ✓ Includes Drivers & Apps.

GET YOUR EMBEDDED CONTROLLER
PROJECT RUNNING FAST!
WITH THESE ACCESSORIES

Relay I/O

\$139
Qty 1 Price



- ✓ 8 10 Amp Relays
- ✓ 8 Opto-Isolated Inputs

JK microsystems
Cost Effective Controllers for Industry

TO ORDER (510) 2364151

FAX (510) 236-2999—email: jkmicro@dsp.com
Visit our WEB site—www.dsp.com/jkmicro
1275 Yuba Ave., San Pablo, CA 94806

The Global Positioning System

Part 2: Implementation Problems and Solutions

GPS receivers tell you how far you are from the center of the earth, not where you are on the earth's surface. Do-While shows us how to get Cartesian coordinates from the geocentric coordinates used by GPS.

FEATURE ARTICLE

Do-While Jones



In Part 1, we looked at how the Global Positioning System (GPS) works.

This month, I'll show you how to use GPS data in practical applications.

GPS data comes in a geocentric (earth-centered, earth-fixed) coordinate system. You usually need to convert the position data to geodetic (earth-surface) coordinates (i.e., longitude, latitude, and altitude) or distances East, North, and Up relative to a particular longitude, latitude, and altitude.

Although ground-based systems use the East, North, Up coordinate system, aircraft find it more convenient to use the North, East, Down system. I'll leave this simple conversion for you.

In addition to talking about coordinate transformations, I'll also discuss integrating GPS data with other navigational data. In particular, I look at a way to combine GPS data with Inertial Navigation System (INS) data to obtain more precise position measurements than either system obtains alone.

GEOCENTRIC COORDINATES

GPS uses geocentric coordinates. As you can see in Figure 1, the x-axis goes from the center of the earth through the equator at 0° longitude near Nigeria. The y-axis goes from the center of the earth through the equator at 90° east longitude near Indonesia. The z-axis goes from the center of the Earth through the North Pole.

This system is wonderful for tracking satellites, but it isn't very convenient for Earthlings. We like coordinate systems that tell us where we are on the surface of the earth, not how far we are from the center of it.

GEODETTIC COORDINATES

Several different earth models are commonly used to express position in terms of latitude, longitude, and altitude. The current favorite seems to be WGS-84 (i.e., the World Geological Survey model of 1984).

The WGS-84 model treats the world as an ellipsoid with a polar radius of 6,356,752.3142 m and an equatorial radius of 6,378,137.0 m. The surface of this ellipsoid is roughly equivalent to sea level. Real sea level differs from the model surface by tens of meters in some places because sea level isn't quite as smooth as you might expect.

Sea level is defined in terms of gravitational potential, not distance from the center of the earth. There are places inside the Earth that are denser than other places, so the gravitational force is greater there. Surveyors use a gravity meter to measure altitude. (The force of gravity can be determined by measuring the period of pendulum.)

There are published tables giving the difference between the WGS-84 ellipsoid surface and sea level as a function of latitude and longitude. Some are classified, but the Defense Mapping Agency technical report DMA TR 8350.2 is approved for public release with unlimited distribution.

In the U.S., the difference varies from -20 to -40 m. If you use -30 m in the equations, then your altitude calculations are accurate to within 30 just about anywhere in the U.S. If you need better accuracy or you're calculating locations outside the U.S., get the DMA technical report.

Altitude is expressed in survey feet. A survey foot is 12.000024". It's so close to a British foot that the difference doesn't matter in most cases.

But, converting distances as large as the radius of the earth isn't one of those cases. The difference is more than 40', so you must use survey feet in calculations involving distances from the center of the earth. By definition, 1200 m exactly equals 3937 survey feet.

Listing 1a shows the constants you need to use for the calculations. Latitude and longitude are normally in degrees, minutes, and seconds. But, computers calculate angles in radians.

Listing 1b shows the equation you need to express angles in radians.

Once you convert the latitude and longitude to radians, you can compute geocentric coordinates (i.e., x, y, and z). First, calculate the values needed in subsequent equations (see Listing 1c).

Determine the radius of curvature (RO) of the WGS-84 ellipsoid at this latitude, as in Listing 1d. Find the height above the ellipsoid in meters using the equation in Listing 1e. Finally, Listing 1f shows all the terms you need to find x, y, and z (in meters).

Of course, you're probably more interested in converting x, y, and z to latitude, longitude, and altitude. But, if you use simple algebra to rearrange the equations in Listing 1, you'll run into trouble.

You need to know RO to find latitude and you need to know latitude to compute RO. Therefore, you have to use an iterative technique to guess the latitude and see if you get the right answer. Keep refining your guesses until you converge on a solution.

I work in southern California at about 35.5° latitude, so I use 35.5° as the initial guess. My first guess is always pretty close, so it only takes a few iterations to converge on a solution. Here's the algorithm for computing latitude, longitude, and altitude.

First, find the longitude (in radians):

$$\text{LONGITUDE} = \text{Arctangent}(Y, X);$$

Convert LONGITUDE from radians to degrees by dividing by DEGREES_TO_

RADIANS. Then, find the minutes by multiplying the fractional part by 60. Multiply the fractional part of minutes by 60 to get seconds.

Second, find the distance from the polar axis:

$$D = \text{Square Root}(X * X + Y * Y)$$

Third, iteratively solve for latitude and height as shown in Listing 2.

Fourth, convert LATITUDE to degrees, minutes, and seconds using the method described for converting LONGITUDE to the same.

Finally, find ALTITUDE from HEIGHT by adjusting for local sea level.

$$\text{ALTITUDE} = (\text{HEIGHT} - \text{LOCAL_SEA_LEVEL}) / \text{SURVEY_FEET_TO_METERS}$$

CARTESIAN COORDINATES

Despite appearances to the contrary, the earth is not flat. (You no doubt knew that already.)

Figure 2 illustrates how the earth's curvature makes noticeable differences in the concepts of East, North, and Up. These differences are obvious for places that are far away. Look at a globe, and compare North from an Alaskan's point of view to North as measured in Greenland.

But, what most people don't realize is that differences in the cardinal directions are significant for points differing from each other by just a few miles. I process data taken from radars separated by 2–10 mi. If I didn't adjust for differences in local North (and East and Up) directions, the errors would be large enough to render the data useless.

The conversion from geocentric GPS coordinates to local East, North, Up Cartesian coordinates can be done easily using a translation vector and a rotation matrix. The trick is figuring out which vector and matrix to use. Since I'm not in the Magicians' Union, I'll reveal the trick to you.

First, pick the latitude, longitude, and altitude of the point you want to use as the origin for your local Cartesian coordinate system: Compute its x, y, z position, and call it ORIGIN.

ORIGIN is the translation vector. To find the position of a GPS receiver

Listing 1—To convert geodetic positions into geocentric coordinates takes several steps. a & b—Some values (a) must be assigned to constants (b) that simplify the calculations. c—Store trigonometric calculations as variables to improve efficiency. d & e—Next, find the radius of the Earth's curvature at your latitude from the equatorial radius and eccentricity of the spheroid (d) and the height above this surface due to local variations in sea level (e). f—Plug in these values to obtain the geocentric coordinates x, y, and z from the geodetic position.

a)

$$\text{DEGREES_TO_RADIANS} = 0.01745329252$$

$$\text{SURVEY_FEET_TO_METERS} = 1200 / 3937$$

$$A = 6378137.0 \quad (\text{A is equatorial radius})$$

$$B = 6356752.3142 \quad (\text{B is polar radius})$$

$$\text{ONE_MINUS_E_SQUARED} = (B/A) * (B/A)$$

$$\text{E_SQUARED} = 1 - \text{ONE_MINUS_E_SQUARED}$$

$$E = \text{Square Root}(E_SQUARED) \quad (\text{E is eccentricity of WGS-84 spheroid})$$

$$\text{LOCAL_SEA_LEVEL} = -30$$

b)

$$\text{RADIANS} = (\text{DEGREES} + \text{MINUTES}/60 + \text{SECONDS}/3600) * \text{DEGREES_TO_RADIANS}$$

c)

$$\text{SIN_LAT} = \text{Sin}(\text{LATITUDE})$$

$$\text{COS_LAT} = \text{Cos}(\text{LATITUDE})$$

$$\text{SIN_LONG} = \text{Sin}(\text{LONGITUDE})$$

$$\text{COS_LONG} = \text{Cos}(\text{LONGITUDE})$$

d)

$$\text{RO} = A / \text{Square Root}(1 - (E_SQUARED * \text{SIN_LAT} * \text{SIN_LAT}))$$

e)

$$\text{HEIGHT} = \text{ALTITUDE} * \text{SURVEY_FEET_TO_METERS} + \text{LOCAL_SEA_LEVEL}$$

f)

$$X = \text{COS_LONG} * (\text{RO} + \text{HEIGHT}) * \text{COS_LAT}$$

$$Y = \text{SIN_LONG} * (\text{RO} + \text{HEIGHT}) * \text{COS_LAT}$$

$$Z = \text{SIN_LAT} * (\text{HEIGHT} + \text{ONE_MINUS_E_SQ} * \text{RO})$$

relative to 0 RI G I N, subtract the translation vector from the GPS receiver's x, y, z position. This isn't quite as simple as it sounds because, in most cases, the x, y, z values of 0 RI G I N and the GPS position are nearly identical.

Remember, the radius of the earth is nearly 4000 mi. If you're going to find a 1-mi. difference between two vectors 4000 mi. long, you'd better be sure you're using adequate resolution.

I like to express geocentric distances using vectors of 32-bit integers with the LSB scaled to one-sixteenth of a meter. This gives a resolution of 2.46" and a maximum range of 166,798 mi. (42 times the radius of the earth or 70% of the distance to the moon).

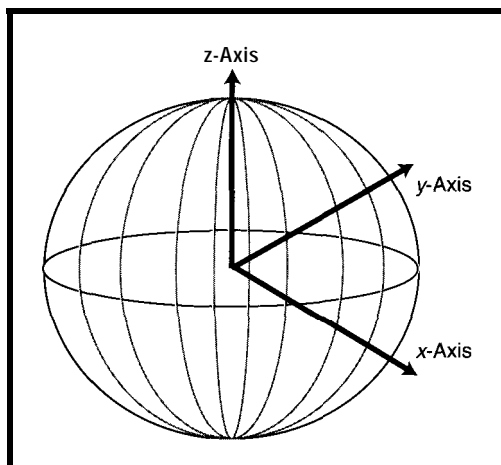
I use integers because finding small differences between large numbers is best done using integers. Floating-point subtraction has some well-known problems with small differences between large numbers unless you use double-precision arithmetic. Besides, integer subtraction is faster.

The difference between 0 R I G I N and the GPS position is a vector based at 0 R I G I N with components in the x, y, and z directions. You probably want a vector with components in the East, North, and Up directions. So, multiply the vector by a rotation matrix specially computed for 0 R I G I N. Here's how you construct that matrix.

Find the x, y, z position of a point above the origin. Use the same values of latitude and longitude as you used for 0 R I G I N, but let the altitude be much higher.

I simply add 3.3 million feet to the altitude of 0 R I G I N to get a new point. Adding 3.3 million feet gives me eight digits of accuracy when I convert to sixteenths of a meter. Eventually, I round the matrix to six digits of accuracy, but I use double-precision floating point

Figure 1--The geocentric coordinate system is an earth-centered, earth-fixed coordinate system. The x-axis goes from the center of the earth through 0° latitude, 0° longitude. The y-axis goes from the center of the earth through 0° latitude, 90° east longitude. The z-axis goes from the center of the earth through the North Pole.



Listing P-Geodetic position can be computed using this iterative algorithm

```

LAST_LAT = 35.5 * DEGREES-TO-RADIANS (or latitude of your choice)
BEGIN LOOP
  SIN_LAT = Sin(LAST_LAT)
  RO = A / Square Root(1 - (E_SQ * SIN_LAT * SIN_LAT))
  HEIGHT = (Z/SIN_LAT) - ONE_MINUS_E_SQUARED * RO
  LATITUDE = Arctanget(
    (RO + HEIGHT) * Z,
    (RO * ONE_MINUS_E_SQUARED + HEIGHT) * D)
  EXIT IF ABSOLUTE VALUE(LATITUDE - LAST LAT) IS SMALL ENOUGH
END LOOP

```

and carry at least two guard digits when calculating the matrix.

Call the vector from 0 RI G I N to this new point ABOVE. The second step is to compute the position of a point about 3 million feet above the North Pole. I do this by changing the latitude of ABOVE to 89°, leaving the longitude and altitude the same.

Subtract 0 RI G I N from the position of this point above the North Pole and call that vector N. N is nearly North, but not exactly because it isn't exactly horizontal, but it doesn't really matter.

What matters is that N is a big vector that has absolutely no East component. In other words, the ABOVE vector and the N vector both lie in the North/Up plane. Take the cross product of ABOVE and N and you obtain a big vector perpendicular to that plane. Call this vector EAST because it is exactly in the East direction.

The third step is to take the cross product of ABOVE with EAST to get a vector NORTH that points in the North direction. Normalize these three vectors by dividing each one by its own length to get the unit vectors in the local East, North, and Up directions.

These three unit vectors are placed in the rotation matrix.

The East unit vector is the first column of the rotation matrix. The North unit vector is the second column of the rotation matrix. The Up unit vector is the third column.

This looks like a lot of work, but you only need to do it once. After you have computed this rotation matrix, you merely need to multiply the x, y, z vectors by the rotation matrix to get East, North, and Up equivalents.

To get from East, North, and Up coordinates to geocentric coordinates, multiply the East, North, Up vector by the transpose of the rotation matrix and add the translation vector.

Order is important! You must rotate before adding the translation vector when going from geodetic to geocentric. You must subtract the translation vector before rotating when you go from geocentric to geodetic coordinates.

In general, finding the transpose of a matrix can be a pain. Fortunately, rotation matrices have a special property. The inverse of a rotation matrix is the transpose of the rotation matrix.

So, the East unit vector is the first row, the North vector is the second row, and the Up vector is the third row of the inverse rotation matrix.

INERTIAL NAVIGATION

Once you have the GPS data in East, North, Up (or North, East, Down) coordinates, you can use it for inertial navigation. In Part 1, I showed you that GPS accuracy is affected by several factors, particularly refraction, multipath, and the selective availability errors intentionally added to degrade accuracy and make the data unsuitable for hostile weapons.

Surveyors eliminate these errors by taking data over a long time (measured in days or hours) and averaging (low-pass filtering) the data. However, moving vehicles can't average GPS data because they're already in a different place by the time they have enough data to average.

A sailboat trying to find Hawaii doesn't really care if its GPS position is off by 100 m. Hawaii is big enough that 100 m won't really matter.

On the other hand, an airplane landing in heavy fog wants its position to be accurate to at least half the width of the runway, and it needs the position information in a timely manner. Airplanes can't use GPS all by itself because of the short-term errors in it.

That doesn't mean planes can't use GPS. It just means that planes (and missiles) need to use GPS in combination with something else. If differential corrections (see Part 1) are available, the plane can use them to improve the accuracy of the GPS data.

But, you can't be sure there will be a friendly neighborhood FAA radio broadcasting differential corrections everywhere you go. (You can't be sure GPS data will be available everywhere, either.)

Therefore, you need something else. That something else is the inertial navigation system (INS), which measures the accelerations experienced by the vehicle. It integrates the accelerations to get velocities and integrates the velocities to get position.

A properly initialized INS computes position correctly for a while, but integrators tend to drift. As time goes by, position data becomes less and less accurate. Inertial systems suffer from low-frequency errors. Fortunately,

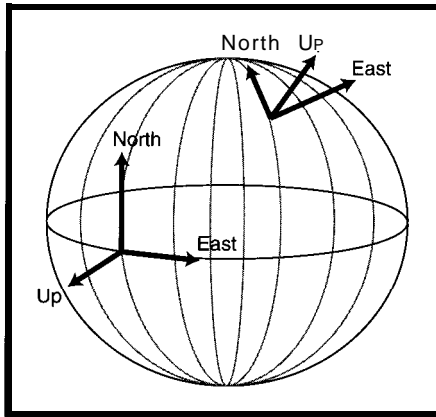


Figure P—Geodetic coordinates have their origin on the surface of the earth. Their notions of East, North, and Up depend on the location of the origin.

there's a clever way to combine GPS and INS data to eliminate all the errors.

Imagine you had two radios tuned to the same station. One is an old tube-type radio with a lot of 60-cycle hum from the filament heaters. The other is a transistor radio with a lot of popcorn noise (above 1 kHz) in one of the stages. How can you combine the outputs from the two radios to eliminate the noise?

Simply adding the two outputs together improves the signal-to-noise ratio by 3 dB because the signal is doubled, but the noise isn't. You can do better than that by adding filters.

Run the tube-radio output through a 200-Hz high-pass filter. Then, use a 200-Hz low-pass filter on the transistor radio. By combining the two filtered outputs, you eliminate the noise.

The problem with this solution is that a signal spectrum in the neighborhood of 200 Hz isn't very flat (unless you're careful about the filter designs). There's a much better way that uses a single filter.

Suppose you subtract the transistor-radio output from the tube-radio out-

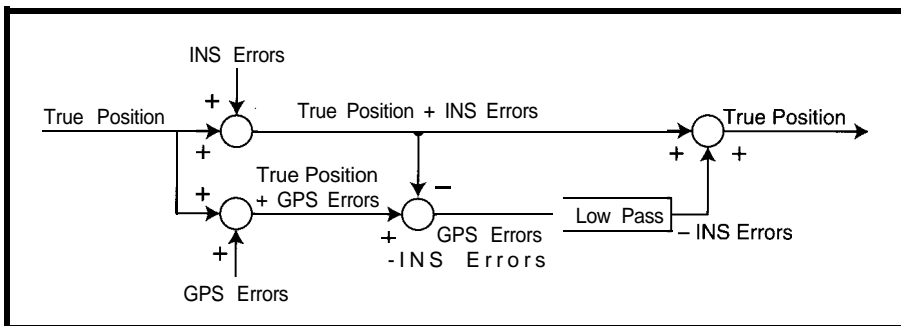


Figure 3—Inertial navigation system (INS) drift is a low-frequency (near DC) drift. GPS errors are higher frequency noise. The scheme illustrated in this block diagram can be used to remove the errors.

CIRCUIT CELLAR®

THE COMPUTER APPLICATIONS JOURNAL

ARTICLE REPRINTS

Circuit Cellar INK owns all copyrights to all articles published in Circuit Cellar INK, the Computer Applications Journal. This copyright also applies to articles downloaded from the Circuit Cellar BBS. Articles may not be reproduced without permission from Circuit Cellar INK.

Authorized reprints of articles are available in minimum quantities of 100. Master copies for duplication up to 500 copies are also available. Call for pricing. All advertisements are removed from articles for reprinting.

All authorized reprints are professionally printed and show the Circuit Cellar INK logo on the first page. Reprints may be printed in black and white or color.

Special permission is available for teachers who would like to copy and use articles in their classes.

For information or for a reprint sample, contact Jeannette Walters at (860) 875-2199 or subscribe@circellar.com.

put. That completely eliminates the signal and leaves just the 60-cycle hum and the popcorn noise. A low-pass filter removes the popcorn noise and leaves only the 60-cycle hum. By subtracting the 60-cycle hum from the tube-radio output, you get just the pure signal.

Since INS errors are low frequency and GPS errors are high frequency, you can use this technique to obtain error-free position information in a timely manner. The GPS position is subtracted from the INS position, yielding the INS drift error and the high-frequency GPS errors.

Put the error data through a low-pass filter to obtain the INS drift error. The INS drift error is subtracted from the INS data to obtain corrected position information (see Figure 3).

Kalman filters are generally used for the low-pass filtering. These filters are fundamentally different from classical analog filters (e.g., Chebychev, Bessel, etc.) as they are digital algorithms that operate on state vectors.

In a typical Kalman filter for a GPS application, the state vector contains

the current acceleration, velocity, and position. The filter then predicts the next position and compares it with the measured position.

It computes the best estimate of the true position from the prediction and measurement and updates its state vector accordingly. One could devote an entire article or book to Kalman filters. Several people already have [see References].

WANT MORE?

In Part 1, I showed you some basic equations to help explain how GPS works. But, these simplified equations aren't sufficient for building a practical GPS receiver.

You have to correct for such things as satellite positions and atmospheric refraction. You have to know the message format and the PN sequences used. There are enough of these necessary details to fill a 137-page document.

The name of that document is "NAVSTAR GPS Space Segment/Navigation User Interfaces," better known as "ICD-GPS-200." It is con-

trolled by the GPS NAVSTAR Joint Program Office SMC/CZ (AFMC) at the Los Angeles Air Force Base.

You can obtain more information about GPS on the Internet. There are too many Web sites to list, but three deserve special mention. A good place to start is Paul Tarr's Web page. Unfortunately, Paul may not be maintaining this site much longer.

The next place is the GPS Joint Program Office home page. Finally, the Coast Guard also has a good site.

These sites have lots of links to other sites to keep you busy surfing for hours. But if you understand GPS, at least you'll be able to figure out where you end up. □

Do- While Jones has been employed in the defense industry since 1971. He has published more than 50 articles in a variety of popular computer magazines and has authored the book Ada in Action. You may reach him at do_while@ridgecrest.ca.us.

REFERENCES

Texts:

S.M. Bozic, *Digital and Kalman Filtering: An Introduction to Discrete-Time Filtering and Optimum Linear Estimation*, Halsted Press, 1994.

R.G. Brown and P.Y. Hwang, *Introduction to Random Signals and Applied Kalman Filtering with Matlab Exercises*, John Wiley & Sons, New York, NY, 1996.

C.K. Chui and G. Chen, *Kalman Filtering. With Real-Time Applications*, Springer-Verlag, New York, NY, 1991.

Web sites:

Coast Guard, <http://www.navcen.uscg.mil/>

GPS Joint Program Office, <http://www.laafb.af.mil/CZ/homepage/>

Paul Tarr's Web page, http://www.inmet.com/~pwt/gps_gen.htm

IRS

413 Very Useful

414 Moderately Useful

415 Not Useful

Serious equipment for serious weather watchers

Accurate, state-of-the-art, and affordable, our weather stations let you monitor the weather like a professional!



Weather Wizard III™

Monitor indoor and outdoor weather conditions at the touch of a button. The Weather Wizard III offers the most requested weather station features at an affordable price. Only \$195!



Weather Monitor II™

The deluxe personal weather station, combining all our weather monitoring functions into one incredible package. A professional weather station at a price that's equally incredible! Only \$395!



WeatherLink®

Connects your weather station to an IBM or Macintosh for in depth analysis. Compare track, plot, and store data. View real-time conditions on weather bulletin. Only \$165!

Weather Station Features Include:

- Inside & Outside Temperature
- Wind Speed & Direction
- Wind Chill
- Time & Date
- Alarms
- Highs & Lows
- Instant Metric Conversion
- Barometer*
- Inside Humidity*

Optional Accessories

- Rain Collector (\$75)
- Outside Humidity Sensor (\$125)*

* Indicates feature or option for Weather Monitor II only.

Call us for a free catalog

1-800-678-3669

or visit us at

<http://www.davisnet.com>

M-F 7 a.m. to 5:30 p.m. Pacific Time • FAX 1-510-670-0589 • M/C and VISA
One-year warranty • 30-day money-back guarantee • CO373V

DAVIS INSTRUMENTS • 1463 Diablo Ave Hayward CA 94542

50 MicroSeries

68 From the Bench

74 Silicon Update

81 ConneCTime


MICRO SERIES

Steve Ciarcia

Applying the HCS II

Part 2: A New Approach to High-Performance System Expansion

Since its introduction four years ago, HCS II has been a real hit. Here, Steve introduces Answer MAN, a new network module that adds low-cost, low-power analog and digital I/O to the HCS network.



If I asked you to spend \$500 on a home-control system, you'd probably say no. So what if you have to bend over and pick up the TV remote to change channels or scream at the kids to shut off the lights? After all, \$500 is a lot to spend just for convenience.

But, if I asked you to spend \$500 for a top-of-the-line programmable-response security system, you'd be elated to find such a discount. So what if you have to string a few wires or write a few lines of program code—you'd feel like you knew the system better that way. After all, \$500 is nothing when it comes to security.

If this scenario sounds funny, don't laugh. This problem is exactly what's hindering the advance of commercially installed home automation and control systems.

We can justify home control under the name of security but not under the guise of pure convenience. Of course, a well-mannered HCS is both.

I must admit I'm no different. The real reason for my HCS installation is primarily security. The HCS is an automated surveillance and tracking system that warns perpetrators to stay away or, if penetrated, notifies authorities and gathers significant evidence. When not performing guardian chores,

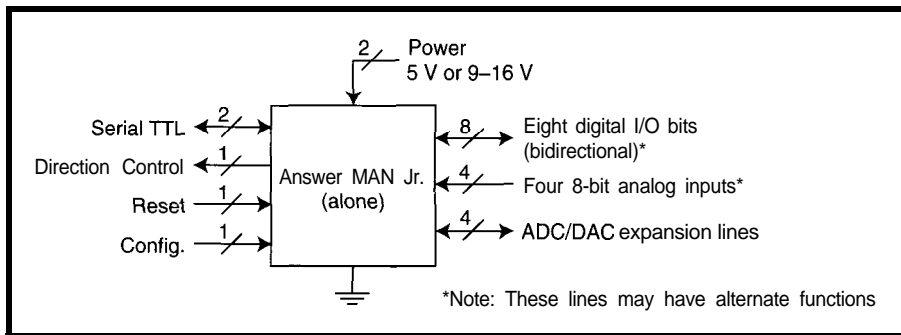


Figure 1—This \$49 module has all the “smarts” to become a universal HCS network expansion module.

the HCS monitors and maintains the proper living environment and supplements the living conveniences.

That’s the long-winded definition of the Circuit Cellar HCS II. For the most part, it’s reality. The HCS II is a specialized embedded control system intended for domestic use.

Anyone installing the HCS II starts with a basic three-board system—supervisory controller (SC), PL-Link, and RBuf-Term. The supervisory controller runs the XPRESS program and coordinates system communication. The SC has analog inputs, a serial RS-485 expansion network, and 24 bits

of digital I/O. The RBuf-Term idiot-proofs the connections to these 24 bits through relays and electrically protected inputs.

The SC or the SC with RBuf-Term is all you need for a direct-wired control system. If you intend to merely monitor a few temperatures and switch on the fans, heating, and air-conditioning components of an HVAC system, you have all the hardware necessary.

Security, on the other hand, usually involves the physical activation of many other electrically controlled devices (e.g., lights, fans, bells, sirens, VCRs, cameras, etc.). Since most of

these devices are AC powered, the typical low-cost remote-control method is X-10. In the HCS, the EL-Link X-10 transceiver enables the HCS to send X-10 control commands as well as listen to the power line.

AN X-10 HEADACHE

Sending and receiving X-10 signals is really where this story starts. The issue is HCS expansion. How far versus how secure? Can you really trust X-10 for more than convenience? Is real security a hardwiring issue?

For the most part, X-10 is reliable. In a small house or apartment, X-10 signals always seem to get through. The occasional missed light when you’re vacuuming is annoying but not traumatic.

Even in security mode, the odds that these same lights and an appliance-module-controlled VCR will activate properly are very high. The odds stay in your favor as long as your neighbor doesn’t mess with his X-10, the AC power is actually on, someone in the neighborhood doesn’t take up arc welding as a hobby, or you don’t

extend X-10 too far. Of course, there’s no real definition of “too far.”

I also use X-10 for convenience lighting control as Ken outlined in Part 1. Lights go on or off based on room occupancy (motion detectors), outside lights come on based on time and light level, and music follows us from room to room. We even use a similar manual X-10 controller to send a bedtime command to the HCS, so the bedroom lights aren’t activated by a restless sleeper.

Like Ken, my HCS installation is designed as a background task-master-not a system where I have to continually negotiate who’s in control. The major difference is that my HCS system is considerably

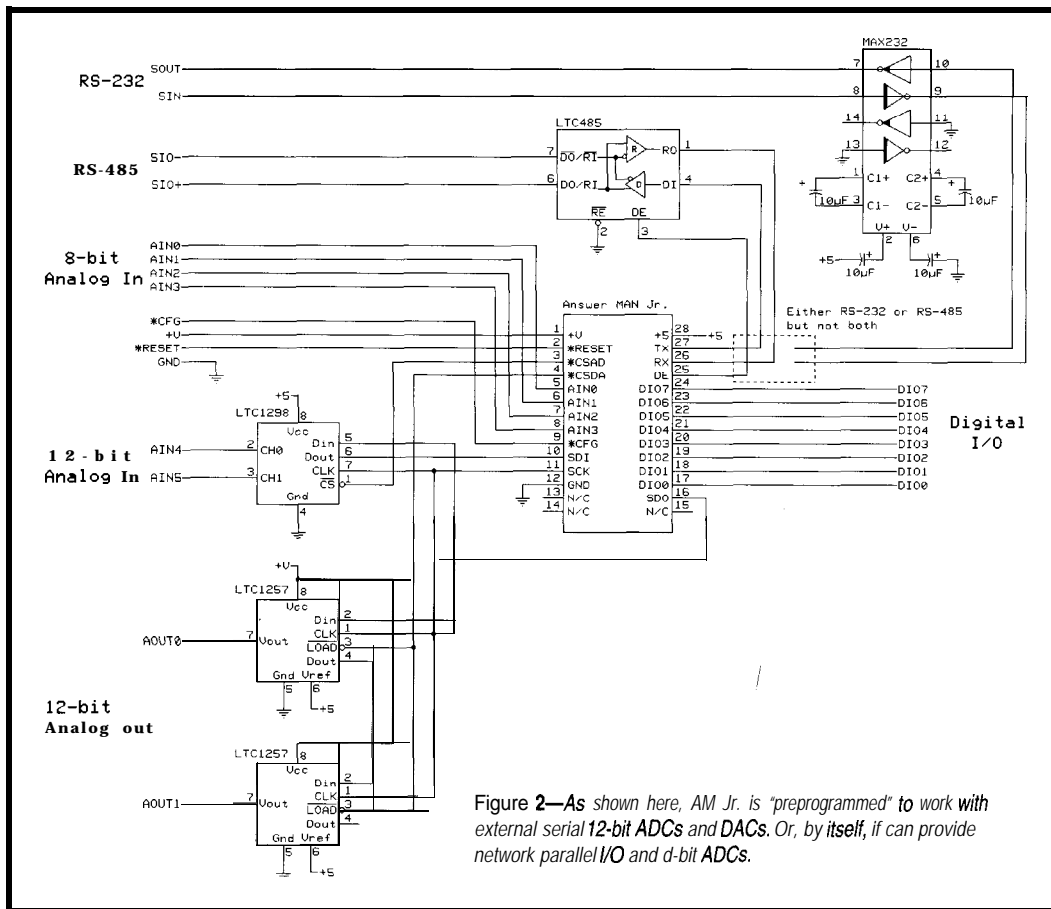


Figure 2—As shown here, AM Jr. is “preprogrammed” to work with external serial 12-bit ADCs and DACs. Or, by itself, it can provide network parallel I/O and d-bit ADCs.

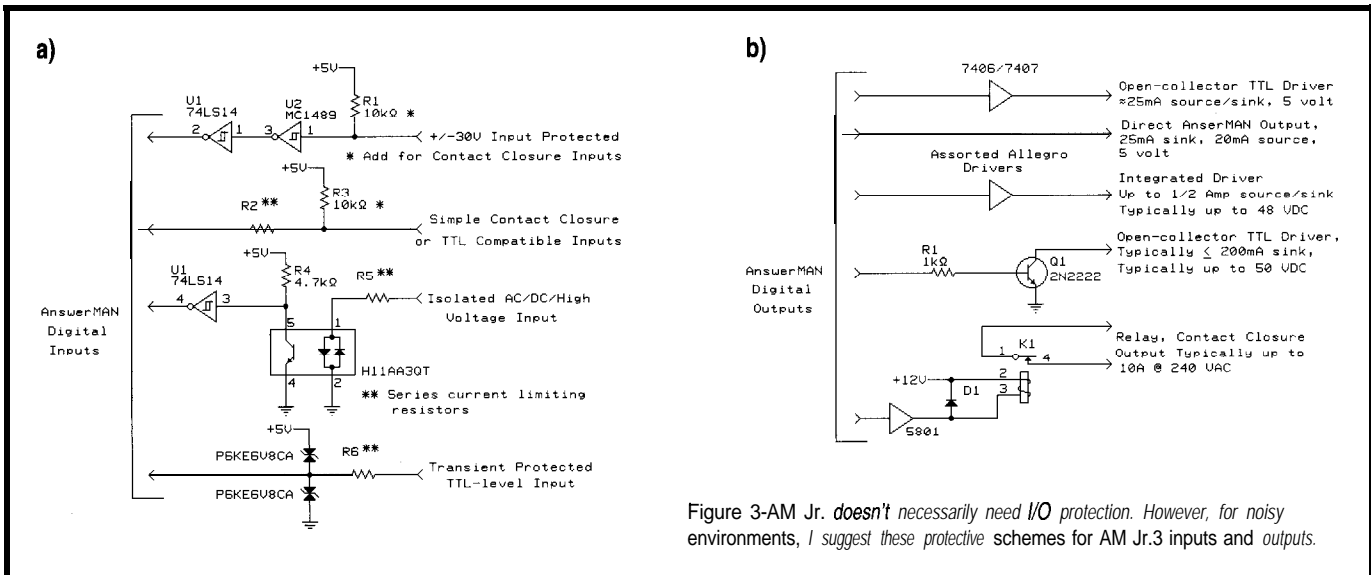


Figure 3-AM Jr. *doesn't* necessarily need I/O protection. However, for noisy environments, I suggest these protective schemes for AM Jr.3 inputs and outputs.

larger because there's more I want to control. That's where trouble begins.

The "real" Circuit Cellar has evolved over the years. While the cellar part still exists with all the electronic junk, there have been so many additions that where I live is now better described as a compound. An HCS system initially adequate for a large house now needs special expansion considerations to cover a larger house surrounded by outbuildings.

The first phase of expansion went smoothly. To monitor garage doors, motion detectors, thermostats, and the like, I ran the RS-485 HCS network out to the buildings and attached DIO-Links. I hardwired solid-state relays for the critical things, but I used X-10 modules for the rest of the AC-powered devices because there were so many of them.

For three years, I've regretted that I didn't hardwire it. These outbuildings are large, and there's a separate 100-A electrical service just for them.

X-10 is supposed to cover everything on the same side of the power transformer, but don't believe it. That only works on odd Thursdays when the weather is above 50°F and cloudy.

The only way to get X-10 working at all was to hardwire a separate TW523 X-10 transmitter from the HCS's PL-Link and install it at the service entrance of the farthest building. You'd

think that, if transmitting into two places helped, transmitting at four or six locations might be even better.

But, when I added a solarium to the house, I had trouble receiving X-10. I added a third TW523, and it seemed to work again. However, when I added a fourth unit (anywhere), it crashed the whole X-10.

I'm no RF expert, but apparently, it was just coincidence that the placement of my three transmitters worked. Each additional transmitter introduced phasing problems and line ringing in the various transmission line lengths. It's a wonder it worked at all. I don't recommend this practice to anyone.

THE GREAT MELTDOWN

One day, one of the outbuilding garage-door openers died. I had a new one installed, but of course, the new remote wasn't compatible with the

pile of others I already had. To compound the problem, the installer wouldn't have any more for four months.

Undaunted, I plugged in an X-10 RF wireless receiver with an appliance module controlling a relay. The relay contacts were wired across the manual door control.

I pressed the hand-held transmitter from the truck and presto-the door went up. Pressed it again and it went down-just like a regular remote. Not only was the problem solved, but because the HCS "hears" X-10 commands, it also knew who opened the garage.

Well, that worked for a few weeks. And then, my wife simply drove in the driveway. The lights went on as usual. Then, X-10-controlled high-intensity strobe lights activated along with some of the outside lights, and a beeper started. And, worst of all, that one garage door opened.

A classic case of total X-10 control meltdown!

I have only myself to blame. I broke the cardinal rule of "serious" home control. Don't use X-10 for anything important! A garage door opening by itself constitutes a complete security failure!

BACK TO THE DRAWING BOARD

I still had the basic problem of extending control to these

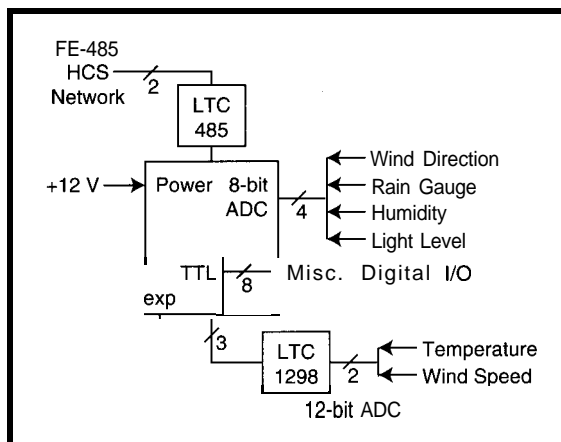


Figure 4-The six analog weather inputs connect to AM Jr. as shown.

Photo 1—About the only way to calibrate an anemometer is to apply some wind. I did it by driving down the highway with it propped against my sideview mirror.



outbuildings. If I considered X-10 undependable, the only choice was to use the HCS RS-485 expansion network and put more link modules all over the place.

However, the total number of control points could easily push the 32-module limit given the present application-specific net-links. My system currently uses eight network modules (i.e., three DIO-Links, PL-Link, DIO+-Link, MCIR-Link, and two LCD-Links). Adding another keypad, reading a dozen analog values, controlling a dozen lights, running an LCD display, and monitoring a pile of sensor contacts would take a lot of network modules.

The solution was as big as the problem. The HCS II was designed as a digital control device, not a remote data-acquisition system. There are eight analog inputs on the SC board itself, but the HCS II doesn't accommodate a lot of external analog monitoring. In the years since its inception, however, we've come to see the need for more analog sensing.

ANSWER MAN TO THE RESCUE

As soon as that garage door opened by itself, making a better and more cost-effective networked I/O system on the HCS became my only goal. What I ultimately needed to solve the present problem, as well as provide for the future, was an RS-485 network link incorporating a bunch of digital and analog I/O in the same package.

I'm embarrassed to say the solution was right under my nose in the Micromint shipping department. Because the HCS has always been a Circuit Cellar-only product, shopping for a solution wasn't Ken's or my first thought.

However, the cost-effectiveness of adapting an off-the-shelf device like Answer MAN quickly justified its adoption. It just remained for Ken to

adjust the HCS's XPRESS protocol to treat Answer MAN as just another HCS network peripheral.

Answer MAN is a multidrop network-based data-acquisition and control module most typically used as an interface between sensors and a computer.

It comes in a 28-pin enclosed DIP package (all the specs are detailed on the Web at <<http://www.micromint.com/>>). Because it responds to ASCII Set and Query commands, similar to other HCS network modules, it's easily adapted to XPRESS.

Although Answer MAN comes in two versions, Answer MAN Junior is the one I'm referencing for HCS use here. At \$49, it's extremely cost-effective and easily expandable for HCS-specific applications.

The Senior version includes all of AM Jr's functions plus the hardware drivers, ADC, and DACs that can be attached to a Junior as expansions. My applications all use AM Jr. with usu-

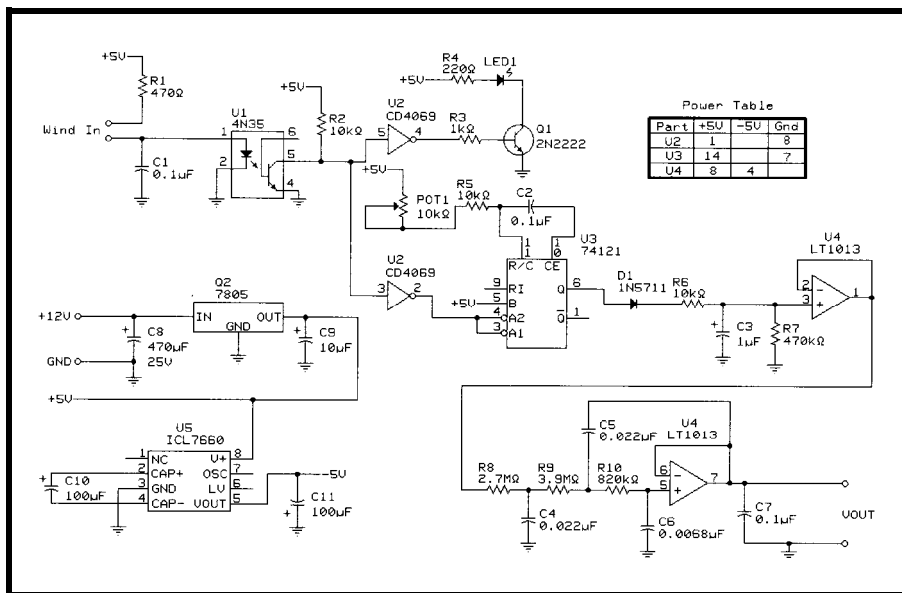



Figure 5—This circuit integrates the anemometer-pulse output into a reasonably steady DC voltage level indicating speed in miles per hour.

Put Circuit Cellar INK to work for you

Circuit Cellar INK readers are engineers, programmers, consultants, and serious computer technologists. Bring them into your store by selling **Circuit Cellar INK**. **Circuit Cellar INK's** Direct Dealer Sales Program lets you increase your store traffic—**increase your sales—**increase your profits with no risk. Let our advertising efforts turn into your profits. For information on how your business can become part of the growing **Circuit Cellar INK** success story, contact us at:

Circuit Cellar INK, attn: Dealer Sales
 4 Park Street • Vernon, CT 06066
 (860) 875-2199 • Fax (860) 871-0411

Figure 6—Using the anemometer frequency counts at the various speeds, I was able to graph its pulse output versus speed.



ally only one or two extra chips. If you decide on AM Sr., note there is a pinout difference between it and Junior.

Figure 1 presents a block diagram of the AM Jr., and Figure 2 schematically demonstrates its full expansion capabilities. A development board is available that allows easy addition of these chips to test your application.

That said, using AM Jr. is much like any other HCS network module. Connect the module directly to a PC running a terminal emulation program (e.g., Procomm), ground the configuration pin, and turn on the power. It comes up at 9600 bps in configuration mode and presents a sign-on banner.

After you remove the configuration jumper, simply follow the directions to configure the module for the operations you want. First, designate the name of the module. For the HCS, 16 modules are allowed with the designations MAN0–MANF.

Next are the module options. While Answer MAN supports many kinds of I/O options, only certain combinations can be used simultaneously. Also, in the case of AM Jr., using the 12-bit ADC and DAC functions requires the chip's external connection.

There is a 16-bit configuration flag register. AM Jr. is smart enough not to allow incompatible configurations. When in doubt, just set what you think you'd like as a configuration and then read back the configuration register. The valid compatible options remain set.

Power down the module, and the configuration is complete. The ID and configuration are automatically saved to EEPROM.

The next time you power up AM Jr., it comes up operating as you configured it. The only job left is to string an RS-485 line out to the area in question, attach AM Jr., and write a little XPRESS code.




Photo 2—Most of the weather interface was mounted on a metal plate and installed in a protective box. The AM Jr. network interface itself is located in the upper left.




Figure 7—The easiest way to convert the wind-direction-resistance output into a measurable voltage is to apply a constant current

INTERFACING TO THE REAL WORLD

OK, so adding HCS control to an outside building is now as simple as running a twisted-pair RS-485 and screwing on an AM Jr. With all that I/O capability, monitoring three garage doors, two motion detectors, two water sensors, and two temperatures, along with controlling three outside lights, an automatic lock, and a couple HVAC fans is simply a matter of wiring everything to Answer MAN, right? Well, yes and no.

The fact that you're reading **INK** already establishes a certain level of expertise, so bear with me. For those people who can write slick XPRESS programs but still put in screws with a hammer, I'd like to point out that there are various ways to connect AM Jr. to sensors and controllable devices. Exactly how is more often a measure of experience than any hard and fast electrical rules.

Figure 3 demonstrates various input and output protection schemes for digital I/O such as Answer MAN. How much, if any, you incorporate depends mostly on the application.

If you're operating low-current devices like LEDs, you can drive them directly or by adding a single transistor. When dealing with motors, fans, or high-current devices, use AM Jr.'s output to control a large mechanical or solid-state relay and route the high current through the relay's contacts.

Digital inputs are also a consideration. For the most part, damage is usually the result of too much input current (i.e., from too much applied voltage). Sources might be a 12-V logic level from a motion detector or high-

EMBEDDED BIOS™

BIOS Adaptation Kit

Just a few of our satisfied customers!

"We're impressed by the level of documentation and particularly by the readability of the code"
- M. Ryan, Cordant, Inc.

"Fast, flexible, high-quality code, and excellent technical support."
- L. Allen, Andros, Inc.

"Personally, I found the Adaptation Kit and tool set very straightforward to use, making the BIOS development process relatively easy."
- S. Chaplin, Software Engineer.

Step up to Embedded BIOS

- Includes our award-winning Run-From-Rom DOS ●
- Includes Flash Disk for popular Flash parts ●
- Over 300 easy configuration options ●
- BIOS Kit includes full source code ●
- Royalties **-\$4/copy & down** ●

Phoenix, Award, AMI, SystemSoft

Using one of these BIOS vendors? Switch over to Embedded BIOS for FREE! Regular Price \$4995*

*Proof of existing license. Minimum 1000 units commitment. All TM names are property of their owners. Limited time offer.

For more info., Call 1-800-850-5755

General Software™

320 108th Ave. N.E., Suite 400 • Bellevue, WA 98004
Tel: 206.454.5755 • Fax: 206.454.5744 • Sales: 800.850.5755
http://www.gensw.com/general • E-Mail: general@gensw.com
© 1996 General Software, Inc. General Software, the GS Logo, Embedded BIOS and Embedded DOS is a trademark of General Software. All rights reserved.

FREE-DEMO

#120

It All Adds Up!

- + A/D inputs with 12-bit accuracy
- + Analog outputs
- + Intel and Ziloga CPUs
- + Quadrature encoder inputs
- + Buffered RS-232/485 serial ports
- + Keypad & LCD display interface
- + Program using a PC
- + 512K program, 512K data memory
- + 5V only
- + Built-in BASIC supports hardware
- + Floating point math
- + Starting from \$195

- ✓ Easier, faster development
- ✓ More integrated design
- ✓ More \$\$ savings

With our line of embedded controllers and accessories, everything adds up to make your job a success!

For more information **303-690-1588**
& FREE catalog, call, email or fax: **303-690-1875**

Web Site: www.remotep.com info@remotep.com

REMOTE™
PROCESSING

The embedded control company

#121

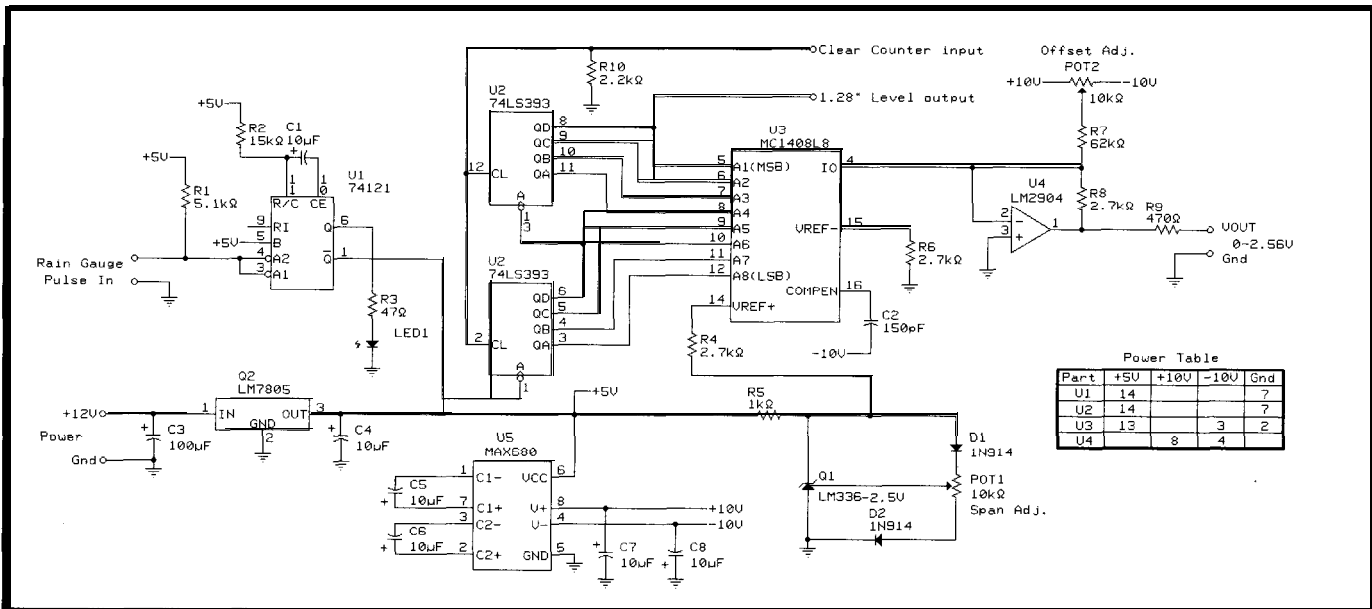


Figure 2 The rain-gauge output is a series of pulses which each represent 0.01" of rainfall. This circuit counts those pulses and converts them to a voltage through a DAC.

voltage transients during thunderstorms.

Recommended protection techniques include series resistors, transorbs, or even optoisolators. One low-cost method I like uses an MC1489 RS-232 receiver as a wide-range DC input device. Capable of withstanding ± 30 -V levels, the 1489 still switches at 0-5-V TTL inputs as well as 12 or 24 V.

ADDING A WEATHER STATION

Attaching AM Jr., with its high-resolution analog inputs and intelligent digital functions, opens up a whole lot of HCS network applications that were cost-prohibitive before. While my first priority was augmenting outbuilding security, it took about 3 ns to realize that AM Jr. was also the perfect interface for adding a weather station to the HCS. And, since Answer MAN talks serial ASCII, this whole weather station can be attached to any RS-232/-485 serial line.

I don't mean just reading a temperature or two. I mean collecting weather data like wind speed and direction, humidity, temperature, rainfall, and light level.

Once acquired by the HCS, this information is formatted and displayed on a 4 x 20 LCD in the kitchen. As data is acquired, the HCS can shut off

sprinklers and retract awnings when it rains, turn on lights at dusk (independent of time), report local weather conditions when you call your HCS while on vacation, or record weather conditions versus energy consumption on a regular basis.

Of course, the last time I made a remote weather station was in February 1982's *BYTE*. Back then, I adapted a wind-direction indicator and anemometer from Heathkit. The anemometer was pulse output, as you might expect, and the direction indicator used a four-bit Grey-code position indicator with optical sensors.

This time, the predominant source of home-based weather sensors seems

to be Davis Instruments. Davis makes a complete line of weather sensors and indicators which are intended to be used with their electronics, of course.

The good news is that most of the sensors are available separately. I ordered an anemometer/wind-direction indicator and a rain gauge. (For the record, when you call Davis, tell them you're doing the Circuit Cellar project. It might save a lot of explanation about buying pieces of their system without their controller.)

Finally, I swiped a humidity sensor from one of Jeff Bachiochi's projects—no use reinventing the wheel—and made my own temperature sensor. For light level, I glued a solar cell on the side of the box.

Acquiring all the sensors was the easy part. Interfacing to them in the most logical way is the hard task. The anemometer pulses as it turns. Its frequency is proportional to wind speed. Unlike the Heathkit device, the Davis wind-direction indicator has a 0-28-k Ω output which represents 0360" position.

The Davis rain gauge is a tilting-cup positive-displacement volume counter. When 0.01" of rain falls into one side of the cup, it tilts to the other side. As it flops over, a magnet closes a reed switch.





Photo 3—The Davis Instruments rain gauge uses a tilting-cup positive-displacement volume sensor. Every 0.01" of rain tilts the cup.

Serious Test Instruments

500MHz logic Analyzers





Starting at \$1350

Get the **speed** you need with our instruments. Our 500 MHz PC based logic analyzers with up to 160 channels, 512K of memory per channel, and optional digital pattern generation.

200MS/s Digital Oscilloscope & Logic Analyzer

Starting at \$1799




Our 200 MS/s DSO gives you 2 or 4 scope channels with long 128K memory buffers, 8 or 16

channels of logic analysis, and FFT spectrum analysis, all integrated into one card. Easy hardware setup and straight forward DOS and Windows software get you started right away.

Model 3100 Device Programmer

\$475 - Our device programmer PLDs, programmable logic, GAL, EPROM, Flash, serial memories, and microcontrollers.



(201)

369 Passaic Ave • Suite 100 • Fairfield, NJ 07004 • Fax (201) 808-8786

Web: www.LinkInstruments.com • Email: Sales@LinkInstruments.com

Total rain is the number of times the cup fills (and tilts) multiplied by 0.01". While the output could be measured as frequency (in a monsoon), the rain-gauge output is really one for an event totalizer.

Humidity would normally be frequency. The sensing device is a capacitive element connected as an oscillator time base. Changes in humidity change the frequency.

While I could treat this as a frequency measurement, borrowing Jeff's conditioning interface presented it as 0-5 V representing 0-100% humidity. See Jeff's original project ("Vaporware—Revealing Your Humidity," *INK* 57) for a schematic and design details.

Temperature and light level are both analog quantities. For light level, I used a silicon solar cell. In bright sun, it puts out about 3.5 V. Establishing the night/day division was merely a matter of recording the output reading at the appropriate light level and using that as a setpoint.

There are almost too many options for temperature sensing (e.g., RTDs, thermistors, thermocouples, etc.). The generic Davis Instruments temperature sensor I bought appeared to be a silicon diode.

I used an LM334 in my previous weather station. But this time, I decided on LM34 and LM60 devices. The LM34, which I used primarily for inside temperatures, has an output ("C or °F) that directly correlates to temperature. A reading of 1.047 V is a temperature of 104.7". Similarly, 0.662 V is 66.2".

The problem with using the LM34 for outside temperatures is that it must have an offset voltage to make subzero measurements. While not an insurmountable problem, offsets require an extra calculation step to determine the real measurement. If I have to do calculations, I may as well use something that doesn't require an offset supply.

For outside temperatures, I used an LM60. It has a very linear output from -40 to +125°C that's always a positive value. Because it is so small, I soldered it directly to the lead wires and dipped it in protective coating.

That's the whole weather station. Totaling up the requirements, there's

Listing I-Here's the Xpress program to read the weather sensors attached to AM Jr. and to display the results on a 4 x 20 LCD.

```

Define Wind-Speed = Variable(10)
Define Peak-Wind = Variable(11)
Define Wind&Direction = Variable(12)
Define Temperature = Variable(13)
Define Humidity = Variable(14)
Define Rain = Variable(15)
Define Light-Level = Variable(16)

Define humid = ADC(16)
Define wdir = ADC(17)
Define rgauge = ADC(18)
Define light = ADC(19)
Define temp = ADC(20)
Define wspeed = ADC(21)

Define Update-Timer = Timer(0)

Begin

IF Reset THEN
  ClearVariables
  Update-Timer = 0N
  LCD(0) = "\e[2J\e[b\e[7h"
END

IF Update_Timer>=15 THEN
  ! Convert 0-2 V (0-2048) into 0-99 MPH
  Wind-Speed = (wspeed * 10) / 205
  IF Wind_Speed>99 THEN
    Wind-Speed = 99
  END
  IF Wind_Speed>Peak_Wind THEN
    Peak-Wind = Wind-Speed
  END

  ! Convert 0-3.59 V (0-184) into 0-360"
  Wind-Direction = ((wdir / 2) * 195) / 50
  IF Wind_Direction>359 THEN
    Wind-Direction = 359
  END

  ! Convert 174-1049 mV (142-859) into -40 to +100°F
  Temperature = ((temp - 142) * 20) / 102 - 40

  ! Convert 0-5 V (0-255) into 0-100% humidity
  Humidity = (humid * 100) / 255
  IF Humidity>99 THEN
    Humidity = 99
  END

  ! Convert 0-2.56 V (0-131) into 0-2.56 inches rain
  Rain = (rgauge * 195) / 100

  ! Convert 0-1.95 V (0-100) into 0-100% light
  Light-Level = light
  IF Light_Level>99 THEN
    Light-Level = 99
  END

  LCD(0) = "\e[1;1HWind %POMPH Peak %PO ", Wind_Speed, Peak_Wind
  LCD(0) = "\e[2;1HDirection %PO\xDF ", Wind_Direction
  LCD(0) = "\e[3;1HTemp %PO\xDFF Hum %PO% ", Temperature, Humidity
  LCD(0) = "\e[4;1HRain %P2\x22 Light %PO\x25 ", Rain

  Update-Timer = 0N
END

```

the frequency for the anemometer, the variable resistance for the wind direction, and zero to a few hundred pulses for the rain gauge. There's also 0-5 V for humidity, 0-3.5 V for light level, and 0.174-1.205 V for temperature. Quite a mix.

All of these quantities are Answer MAN compatible. The 12-bit ADC, with an LSB of 1.2 mV, allows about 0.25" temperature resolution even at these millivolt levels (amplifying the LM60 with a x5 amplifier improves it significantly).

A NOTICE TO OUR READERS:

Circuit Cellar INK® will occasionally provide a listing of subscribers to vendors with offers of substantial interest to our readers. If you would prefer not to be part of this listing, send the mailing label from the front cover of *Circuit Cellar INK®* along with your request to:

Circuit Cellar INK
Subscriber Service
Department
P.O. Box 698
Holmes, PA
19043-9613

Either the 8- or 12-bit ADC can be used for the others. The real problem is the digital stuff.

The good news is that Answer MAN reads frequency and events. Unfortunately, during frequency measurement, it shuts off the analog inputs. The net result is that it would take three Answer MAN Jrs.-two for frequency and one for analog (the digital can be on either).

BECOMING ANALOG FRIENDLY

The solution to making this weather station cost-effective is to use only one module and make all the inputs fit the available I/O!

Coincidentally, there are six sensors and six analog inputs (four 8-bit and two 12-bit). Figure 4 depicts how these six sensors could connect to one AM Jr. if they're all analog.

Converting the rain gauge, anemometer, and wind-direction indicator from digital and resistance outputs is simply a matter of signal conditioning. Figure 5 shows the schematic for converting the anemometer's 0-100-MPH pulse output to a 0-2.0-V analog level.

As the anemometer cups spin, a magnet opens and closes a reed switch. Since I wasn't using the Davis electronics, I had no way initially of knowing how many pulses/s/MPH [i.e., frequency) to expect. As you can see in Photo 1, I solved that question by

Listing 2—*This code reads temperatures and sets or resets the alarm via an iButton keyring*

```

Config MAN_Link      = 2
Config iButton       = 1,0246F20000000098

Define Steve_Code    = 2763
Define Steve_Key     = 1

Define Inside_Temp   = Variable(17)
Define Outside_Temp  = Variable(18)
Define Entry_Code    = Variable(19)
Define Entry_Key     = Variable(20)
Define Alarm_On      = Input(35)

Define itemp         = ADC(28)
Define otemp         = ADC(29)

Begin

IF Input(10)=TRUE THEN
! Convert 0-1000 mV (0-819) into 0-100°F
  Inside_Temp = (itemp * 10) / 81
  Outside_Temp = (otemp * 10) / 81
END

IF Alarm_On=True THEN
  Entry_Code = KeypadNumber(1)
  IF Entry_Code=Steve_Code THEN
    Alarm_On = False
  END
  Entry_Key = iButton(1)
  IF Entry_Key=Steve_Key THEN
    Alarm_On = False
  END
END
END

```

loading the anemometer onto the car with a frequency counter and driving down the road at different velocities.

Unfortunately, I had some problems with wind gusts, but I ran the same course in two directions to compensate. As the graph in Figure 6 shows, my data points weren't all that far afield.

The frequency is surprisingly low—under 30 Hz at 60 MPH. In any case, I obtained a measure of pulses/MPH and converted it to voltage using the circuit in Figure 5 and Photo 2.

The output of the anemometer is fed through an optoisolator. I did this mostly for contact debouncing and to


REAL-WORLD WIRING

The good news about RS-485 networks is that the specification says that network devices can be spread along a 4000' wire, provided it's properly terminated at the end. And, if you string 4000' of RS-485 straight down the side of a highway, it works fine. The bad news is that it's virtually impossible to wire a house as one continuous wire with termination at one end. The guy who came up with this spec never had to wire a house!

The only practical solution to house wiring is a star configuration. But, a star (the wires to net-links radiate individually from a central point) introduces termination problems as well as a natural tendency to have a bunch of stars radiating from one central star.

To alleviate the problem, Circuit Cellar offers the generic RS-485 buffer/repeater shown in Photo i. Each one extends network distance by 4000', so they are a perfect solution for isolating termination problems in a star configuration, providing extra drive for low-gauge wire, and adding transient protection to a network.

Photo i—*This is the assembled RS-485 repeater. The LEDs light to indicate activity and network data direction.*



The 2" x 2" repeater is available as a \$49 kit. It operates on 8-16 V, covers data rates of 300 bps to 115 kbps, and includes both common- and differential-mode transient protection in both directions. Onboard LEDs indicate activity and data direction.

allow driving an LED that tells me that the thing is turning at all.

Each variable-width anemometer pulse is turned into a precise 1-ms pulse width (calibrating the output span is just a matter of adjusting the pulse width). These pulses are then integrated through a simple low-pass filter that produces a varying DC level output.

Finally, to reduce the low-end ripple, the output is filtered through a 2-Hz low-pass filter. The net result is a 0-2-V output that represents 0-100 MPH.

WIND AND RAIN INTERFACE

Davis Instruments' wind-direction indicator is a potentiometer. I measured it as 0 Ω at 0" and 28 k Ω at 359". While there are many ways to convert this to a voltage, the easiest one is $E = IR$.

Transforming 0-28 k Ω into 0-5 V involves passing a constant current through the resistor and measuring the voltage drop across it. A more useful value than 0-5 V for the HCS is 0-3.6 V (to represent 0-360").

A quick calculation shows that a current of 128.5 μ A through a 28-k Ω resistor produces 3.6 V. Figure 7 is a temperature-compensated constant-current source using an LM334 set at

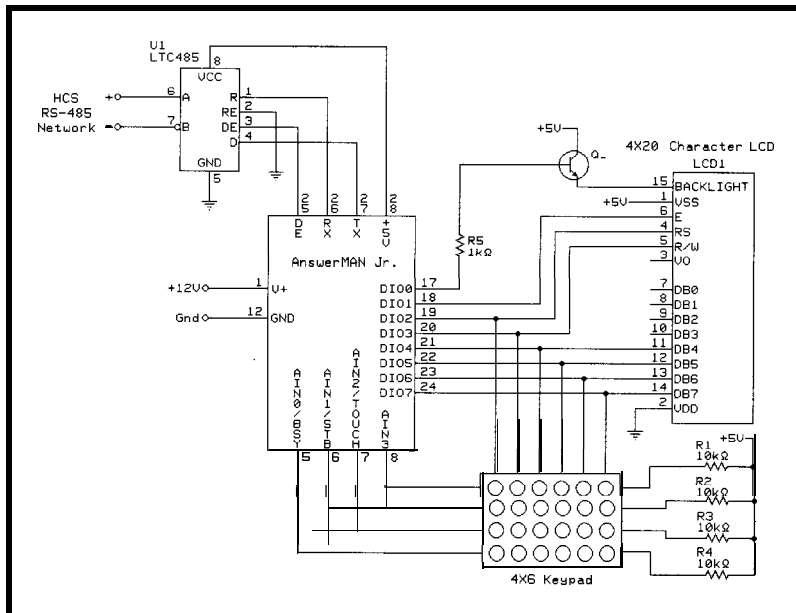


Figure 9-A complete 4 x 20 backlit-controlled LCD with a 4 x 6 keypad makes a low-cost terminal interface using AM Jr. This is equivalent to the former LCD-Link plus a keypad.

128.5 μ A. Even connected to an 8-bit ADC, position resolution is still plus or minus a few degrees, which contrasts with the 23" resolution of my previous station.

Totalling the bucket tilts from the rain gauge and converting it to an analog quantity was equally straightforward. The concept of counting and summing the pulses, as I'd have done if I used Answer MAN's event counter, is valid. The rest is simply digital to analog.

Figure 8 and Photo 3 show the rain-gauge totalizer. The rain-gauge reed switch is first "filtered" through a 50-ms one-shot.

Again, this is mostly for contact debouncing and to flash an LED, so I know it's working (it saves climbing up to the rain gauge). The incoming pulses increment an

8-bit counter connected to an 8-bit DAC.

The rain gauge is calibrated to tilt the bucket (one pulse) for every 0.01" of rain. By adjusting the voltage reference of the DAC to 2.56 V (offset and gain adjustments make it exact), the output voltage exactly represents the total rainfall. A voltage of 1.04 V equals 1.04" of rain.

An 8-bit counter/DAC combination allows up to 2.56". If you live in Florida where it can rain

more in an hour than it might in one month in Connecticut, 2.56" might seem too low.

Therefore, the MSB of the DAC is also connected to Answer MAN as a digital input. When the rain reaches 1.28", this line goes high. It's up to the XPRESS program to clear the counter (a digital output from Answer MAN), save the 1.28" accumulated so far in a variable, and start counting again.

Listing 1 illustrates the XPRESS code for polling the weather station and displaying the collected information on a 4 x 20 LCD. This software illustration is far from complete, and you can do much more. Built-in functions like min and max analog recording allow additional displays of high and low temperatures for the day and so forth without taking up valuable HCS program and variable space.

LOW-COST HCS LCD DISPLAY

AM Jr. has many built-in, high-level features. Among them is a 4 x 20 LCD driver and automatic 4 x 6 keypad scanning.

Building this weather station is a great idea, but having a way to display the data is equally important. Figure 9 shows the complete electronics for an HCS "LCD terminal."

The LCD and keypad connect directly to the module pins. The only extra components are an LTC485 line

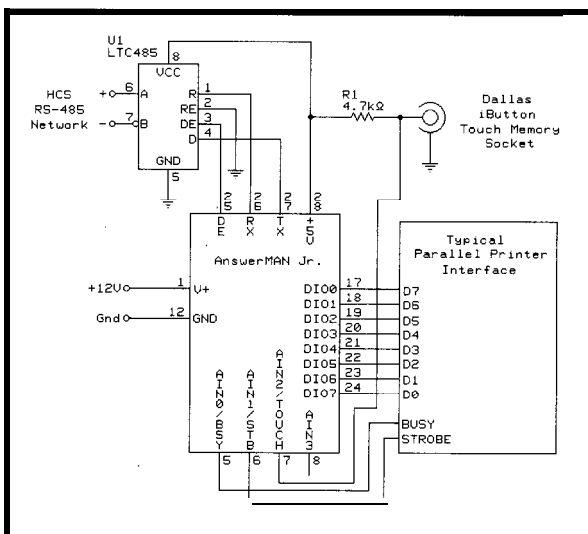


Figure 10—AM Jr. can also be configured to drive a parallel printer or, as shown here, to read the Dallas iButton serial number.

driver and a transistor for back-lighting control. The XPRESS program for reading the keypad and writing to the LCD is in Listing 2. With the exception of more keys, it is programmed the same as the original LCD-Link.

One interesting feature not presented in Figure 9 is that Answer MAN Jr., even fully wired with an LCD and keypad, has additional capabilities. The 2-channel, 12-bit ADC and DAC chips can still be added to monitor and control analog applications.

IDENTIFICATION AND PRINT RECORD

In my HCS, I have a DIO-Link configured to drive a parallel printer. It records the times when the alarm turns on and off and when the driveway or other pertinent sensors are tripped.

Now that I have the weather sensors, I may add printing a daily conditions report to the list. While you can still use the original DIO-Link for this task, AM Jr. has parallel printer output configuration as well (see Figure 10). Again, it really just involves adding an RS-485 driver chip.

To make it more interesting and to demonstrate another new HCS feature for security buffs, note that it only takes a 4.7-k Ω pull-up resistor to add ID protection.

The Dallas iButtons (formerly Touch Memory) have 64-bit serial numbers

Listing 3—Here's the XPRESS code for reading an iButton serial number and verifying it.

```

Config MAN-Link      = 2
Config iButton      = 1,0246F20000000098

Define Steve-Key    = 1

Define Entry-Key    = Variable(20)
Define Alarm_On     = Input(35)

Begin

IF Alarm_On=True THEN
  Entry-Key = iButton(1)
  IF Entry_Key=Steve_Key THEN
    Alarm_On = False
    LPT(2) = "Put the dinner in the oven"
  END
END
END

```

among their prescribed function [the HCS only deals with the serial number]. Each one is factory set, and they're all different.

It's conceivable that you'll put a keypad beside your entry door. But, it's also possible that you could just as easily use iButtons. Give your friends and family iButtons for their key chains or use one of those nice-looking Josten's stainless-steel rings with an iButton on top (see Photo 4). Nothing like a real decoder ring to crank up the HCS.

Listing 3 demonstrates some sample code for driving a parallel printer and reading iButton serial numbers. To make dealing with a lot of 64-bit serial numbers more HCS friendly, rather than trying to compare 64-bit variables

in XPRESS, any serial numbers you wish to acknowledge are individually listed as constants in the XPRESS setup configuration (whether it's 1 or 10 iButtons). When a valid iButton serial number is read, that variable comes back designated as True.

OPENING THE DOOR

The best part about using the new Answer MAN Junior on the HCS network is all the possibilities it opens up. It just so happens that a couple of ABACOM Technology RF transceivers passed across my desk this week. A quick look at the specs suggested half-duplex operation up to 38 kbps.

Let's see, the HCS runs at 9600 bps half duplex. Can I just drop a transceiver in Figure 9 and make a wireless HCS terminal? Stay tuned. Ideas are one thing we have a lot of around here. Be aware, however, that even though we might think of something first, that doesn't mean we want or have the means to capitalize on every nuance.

For example, all the interfaces I described using AM Jr. are hard wired.



Photo 4—Dallas Semiconductor's iButton adds real decoder-ring access to an KS-protected home.

There are no PC boards for AM Jr.'s LCD terminal, weather station, parallel printer, or security interfaces.

Given all the ideas still on the back burner, it may be a long time or we may never make them. With 3000 HCS II systems out there, the opportunity for third-party HCS support, especially for someone making PC boards (with or without supplying Answer MAN), should be obvious.

Remember, this whole place (i.e., Circuit Cellar Inc.) started by sourcing a parts kit for an article. □

Steve Ciarcia is an electronics engineer and computer consultant with experience in process control, digital design, and product development. You may reach him at steve.ciarcia@circellar.com.

SOURCES

Answer MAN Jr.
Micromint, Inc.
4 Park St.
Vernon, CT 06066
(860) 871-6170
Fax: (860) 872-2204
<http://www.micromint.com/>

HCS II, RS-485 Buffer/Repeater Kit
Circuit Cellar, Inc.
4 Park St.
Vernon, CT 06066
(860) 8752751
Fax: (860) 871-0411
<http://www.circellar.com/>

iButton
Dallas Semiconductor, Inc.
4401 S. Beltwood Pkwy.
Dallas, TX 75244-3 129
(972) 371-4448
Fax: (972) 371-3715
<http://www.dalsemi.com/>

#7852 Rain Gauge, #7911 Anemometer
Davis Instruments
3465 Diablo Ave.
Hayward, CA 94545
(5 10) 732-9229

IRS

416 Very Useful
417 Moderately Useful
418 Not Useful

IT'S A NEW YEAR!

SUBSCRIBE TODAY AND SAVE

53% OFF THE DOMESTIC NEWSSTAND PRICE!

Upcoming issues will feature:

February: Distributed Control

March: Home and Building Automation

April: Robotics

May: **Embedded** Processors

PLUS our always-popular BONUS SECTION covering the **EMBEDDED** PC market

One year (12 issues) for only \$21 .95 (U.S.)

\$31.95 Canada & Mexico, \$49.95 all other non-U.S. addresses

(U.S. funds drawn on U.S. banks only)

CIRCUIT CELLAR INK®

THE COMPUTER APPLICATIONS JOURNAL

IT'S EASY TO SUBSCRIBE!

Tel: **(860) 875-2199** • Fax: **(860) 8714411** • BBS: **(860) 871-1988**

Visit our web site at: <http://www.circellar.com/>

Finally, Standard RS-485 Network Software

With Cimetrics' 9-Bit PLAN you can link together up to 250 of the most popular 8- and 16-bit microcontrollers (8051, 80C196, 80C186EB/EC, 68HC11, 68HC16, 68332, PIC16C74).

The 9-Bit μ LAN is:

- ▶ **Fast** — A high speed (62.5k baud) multidrop master/ slave RS-485 network
 - ▶ **Flexible** — Compatible with your microcontrollers
 - ▶ **Reliable** — Robust 16-bit CRC and sequence number error checking
 - ▶ **Efficient** — Low microcontroller resource requirements (uses your chip's built-in serial port)
 - ▶ **Friendly** — Simple-to-use C and assembly language software libraries, with demonstration programs
 - ▶ **Complete** — Includes network software, network monitor, and RS-485 hardware
- Standard-The 9-Bit μ LAN is an asynchronous adaptation of IEEE 1118



e-mail: info@cimetrics.com

55 Temple Place • Boston, MA 02111-1300 • Ph 617.350.7550 • Fx 617.350.7552

Microcontroller
or PC Master

Intel 8051
HVAC/R

Intel
80C186EB/EC

Intel 80C196
Transportation

Microchip
PIC16C74

Motorola 68HC11
Process Control

Motorola 68HC16
Medical

Motorola 68332
Laboratory Data


RS-485 MULTIDROP NETWORK

FROM THE BENCH

Jeff Bachiochi

Liquid Crystal Displays

First You See 'Em, Then You Don't



When you see LCDs, do you think

polarized shades? No? Well, perhaps you should. As Jeff shows, the LCD uses similar technology to accomplish essentially the same thing—reflecting light into just the place you want it.



Who'd want a wristwatch without hands? When it was first released years ago, it seemed like everyone had to have the latest gimmick—an LED digital watch.

Battery replacement became a real drag, however, since LEDs sucked the life from them. The solid-state light just wasn't efficient.

Then along came LCD digital watches. The need to replace batteries fell to an annual nuisance akin to storing away the barbecue grill at the end of summer. The new LCDs didn't create light. They reflected it.

MAGIC WITH MIRRORS

It's all done with mirrors, so to speak. LCDs alter the path of light so it is allowed to or prevented from passing through opposing polarized filters.

Liquid crystals are named for their interesting state between a crystalline solid and an isotropic liquid. Within a particular temperature range, the anisotropic form becomes stable. The crystal-like molecules float in directional patterns like compass needles pointing toward the north magnetic pole.

While their tendency is to align with one another, these crystals remain in a liquid or floating state. They act like tiny lenses, affecting the way light passes through.

All liquid crystals are similar, although the materials used differ in the amount of positional order present. I'll restrict this discussion to one of the simplest mesophases—the nematic phase.

Other than the tendency for these cigar-shaped molecules to align themselves in a parallel fashion, they aren't arranged in any particular pattern or layers, like some of the more complex materials. All liquid-crystal (LC) molecules are affected by electric and magnetic fields.

They may have a permanent polarity associated with them that aids in self-alignment. Like a compass needle, the LC molecule rotates to align itself with an applied current. If the current passes perpendicular to the molecules, they rotate and affect the macroscopic characteristics of the LC material.

Light comes at us circularly polarized, like waves traveling down a jump rope wiggled in all directions. If that rope is placed between slates of a picket fence and wiggled, only the waves that align up and down with the slates filter through. By placing a polarized filter in the path of light, you eliminate waves in all but one direction, creating polarized light.

Liquid crystals are birefringent, which means that they have differing indices of refraction dependent on whether the light is polarized parallel or perpendicular to the molecules. Therefore, if you control the angle that light enters the molecule, you can predict its polarization when it exits.

A change in polarization angle occurs due to the light taking two paths through the liquid crystal—the fastest path (i.e., the ordinary ray) and the slowest path (i.e., the extraordinary ray). When the rays converge, they are out of phase slightly, altering the polarization angle.

So, we have the angle of polarized light that's affected by passing through LC molecules. And, this angle is influenced by the orientation of the light with respect to the molecules' positions.

By applying an electric (or magnetic) field in a way that reorients the molecules, you can alter the amount of polarization change that the light ex-

periences as it travels through the material. But, how does this create a display?

WHAT GOES IN DOESN'T ALWAYS COME OUT

Remember the Polaroid sunglasses commercials? Am I dating myself or what?

This muscular guy at the beach holds up two pairs of sunglasses and rotates the lens of one pair in front of the lens of the second pair. The sun is totally blocked when the two pairs of glasses are at right angles to one another.

The ad demonstrates how polarized light (the kind reflected off the surface of the water) can be blocked by the properties of the new glare-reducing Polaroid sunglasses. Polarized light, whether polarized due to reflection or by passing through a polarized filter, can be passed or blocked by a second polarized filter, depending on its orientation.

Imagine you're looking at two polarized filters positioned so that their polarizations are 90° to one another, with the one closest to your eyes polarized horizontally and the other vertically. Holding them up to the light, just like in the sunglasses ad, the light can't pass through.

If you squirt some liquid crystals (with the proper makeup) between them, as the vertically polarized light from the farther filter passes through the liquid crystals, the angle of polarization changes. If it changes by 90°, it is now horizontally polarized, and it




Figure 1 a-As you can see, liquid crystals aligned by an electric current do not alter the polarization of light. b-However, unaligned crystals change light's polarity.

can pass through the second and closest horizontally polarized filter. Voilà.

DO YOU SEE THE LIGHT?

There are basically two lighting techniques for LCDs—frontlit and backlit. I've just described backlit. The only difference between frontlit and backlit displays is that light travels through a frontlit display twice.

A reflector on the inside sends the light back through the display along the same path it came in on. You may be most familiar with the seven-segment digits on your watch, but the glyphs (segments) can come in any size or shape.

Custom glyph displays are often used on cheap hand-held video games. The more expensive ones use a row and column matrix that any shape can be displayed in.

Each area of the display that needs to be controlled independently must

again pass through the second filter, and the patch appears light.

HOLD THE MAYO

The LCD sandwich is composed of special layer pairs. Working from the outside in, these might consist of polarizers, glass, electrodes, and coating pairs as shown in Figure 2.

The polarizers are set at 90° angles to one another, each on a glass base for rigidity. A transparent layer of metal oxide coats the inner surfaces of the glass. Like circuit traces on a PC board, the patterns form electrodes over the liquid crystals and are brought out to external pins for connecting the necessary driving circuitry.

The two electrode surfaces are coated with a special polymer. The coating is rubbed in the same direction as the polarization of the filters. Because the polarizing filters are 90° to one another, the microgrooves produced from rubbing the polymer are also at right angles.

Liquid crystals are drawn into the cavities between the two layers by a vacuum and the opening is sealed. Once inside, the liquid crystals tend to align themselves with the microgrooves in the polymer coatings.

Since other molecules throughout the LC material self-align with their




Figure 2—The anatomy of an LCD shows the complex layers of material necessary to alter the transmissive properties of liquid crystals.




Photo 1--My prototype came alive thanks to a PicStic processor which has an onboard real-time clock. Push buttons allow the time and date to be set.

This change is accomplished without negative voltage in static (nonmultiplexed) displays by driving the display's common (i.e., backplane) with a clock signal instead of grounding it, as you might with a seven-segment LED.

Segments driven by the same phase clock develop no potential across them, and they remain unaffected or transparent. Segments driven out of phase with the clock develop a plus and minus potential in step with the clock and are tilted into opaqueness.

Since multiplexed LCDs essentially have no common, the polarity reversal is accomplished by creating a fake ground at about $\frac{1}{2} V_{CC}$. This month's project uses a static display. I'll discuss multiplexed displays another time.

STATIC AC DRIVE

The static display I've chosen is a 0.5" 3.5-digit display with decimal points, colon, polarity, and low-battery indicator. It is made by Varitronix and costs about \$5 in single quantities.

It has 29 segment connections plus two backplane commons on a 40-pin, 1.3" center dual-row pin header. A 0.4" four-digit display is also available in the same-size 40-pin package.

Because the AC drive is necessary, it takes a bit of extra circuitry. An

oscillator is necessary to ensure the display continuously receives its AC drive. This can, of course, be created by software when driven by a processor. A circuit failure or a code hangup can ruin the display with a DC bias if it goes

unnoticed for too long, though.

When I looked at shift registers and latched drivers, including the seven-segment LCD driver (4543), the number of parts needed seemed high. While looking through a parts catalog, I found a part I'd overlooked. (Hey, some people look at car or movie magazines. I look at parts catalogs—what can I say?)

Microchip has a 32-segment static LCD driver with built-in oscillator—the AY0438—and it's also about \$5 in single quantities. The chip is driven by a clocked serial shift register with a parallel load control line, so it only requires three output bits. You can get away with two lines if you tie the load line up and clock in the data faster than the segment-switching speed, which is relatively slow.


So, here's a simple and inexpensive LC display with driver for around ten bucks! If you look at the schematic in Figure 4, you'll see a connector where you can drive the display from two or three bits of your favorite circuit. I added a 14-pin SIP socket and a couple push buttons to create a stand-alone LCD clock by adding a PicStic 2 as you see in Photo 1.

The PicStic is a 14-pin SIP microcontroller with 10 digital I/O bits and

BIG THINGS COME

DOMINO

Microcontroller



Starting at \$99

Micromint's Domino-52

microcontroller is a "supercomputer" in less than 0.75 cubic inches. We've packed the most essential elements into one tiny package. Domino is a plug-and-go module, just attach +5V and a terminal or network. A simple keyed sequence saves an autostarting program in nonvolatile memory.

SPECIAL FEATURES

- 80C52 with ROM-resident, full floating-point BASIC
- 32K bytes SRAM and 32K bytes EEPROM
- Two PWM outputs, I²C bus
- Serial I/O: (up to 19,200 bps) RS-422, RS-485 & RS-232A
- Two interrupts and three timers
- Parallel I/O: 12 bits, 3 shared with ADC and I²C
- Power: +5V @ 15 mA;
- Size: 1.75" x 1.062" x 0.4" potted
- A/D converter: 2 channels, 12 bits, 10k samples/sec.
- Connections: via 2-10, 0.1" dual-row header
- -20°C to 75°C operating temperature
- Industrial temperature available



4 Park Street • Vernon, CT 06066

Call 1-800-635-3355

(860) 871-6170 • Fax (860) 872-2204




Sets The Pace

Am5x86™ 133 MHz
PC/104 cpuModules Exceed
Pentium-75 Performance



\$995




The CMV586DX133 offers versatile embedded functionality

ISA Bus Intelligent Data
Acquisition and Control
Boards Multiply Your PC's
Processing Power



\$3985



The IDAC5210 with PC/104 extension bus, 500 KHz analog and digital front-end is powered by on-board Am5x86 133MHz CPU and TMS320C50 DSP

Our PC/104 and ISA Bus product lines feature Analog and Digital I/O, CPU, DSP, Shared Memory, SVGA, PCMCIA, CAN Bus and Intelligent GPS



Real Time Devices USA

200 Innovation Boulevard • P.O. Box 906
State College, PA 16804-0906 USA

Tel: 1 (814) 234-8087 • Fax: 1 (814) 234-5218
FaxBack®: 1 (814) 235-1 260
www.rtdusa.com • E-Mail: sales@rtdusa.com

RTD Europa

RTD Scandinavia

Budapest, Hung&y
Fax: 36-1-326-6737

Helsinki, Finland
Fax: 358-9-346-4539

RTD is a founder of the PC1104 Consortium and the world's leading supplier of intelligent ISA DAS interfaces.




Figure 5-A digit segment can be used to indicate the day of the week by selecting the necessary segments.

a real-time clock/calendar. This miniature controller uses EEPROM program storage and can be easily programmed and reprogrammed in assembly language, C, or PICBASIC.

DIGIT DATA

Microchip's AY0438 driver shifts input data on the rising edge of the clock line. The four bytes of data are shifted in LSBit, starting with the rightmost digit.

When all 32 bits are shifted into position, the load line latches all data into the output latches. The latches can be held in the transparent state by tying the load line high.

If data is clocked in slowly in transparent mode, it shifts through looking like garbage. If clocked rapidly enough, it changes faster than the LCD can respond, so it looks normal as long as you pause before shifting in new data.

I wired the LCD so one byte could be used for each digit and they would, for the most part, all be the same. The first seven bits control the normal segments a-g (see Figure 4).

The last bit is used for the decimal point to the right of each digit. This way, all digits can use the same lookup table. The leftmost digit, which can only be a 1 or blank, handles all the other glyphs on an as-needed basis.

The real-time clock (RTC) passes data in a BCD format, each byte holding a decimal digit in each nibble. Although in the past I might have been forced into writing this in assembler, I now have the PICBASIC compiler at my disposal. I still find BASIC prudent for those times when I don't have hours to spend on software.

SOFT STUFF

Three routines are needed for this project. The first-and by far, the simplest-is the display routine. The display's digits are numbered first through fourth, left to right. Four bytes are needed for the display update.

The digit data is set on the data output line starting with the LSBit of the fourth byte. The clock line is cycled (high and low) after each data bit is set to shift the data bit into the display driver. This process continues for 32 bits until all four digits are sent.

Finally, the load line is cycled (high and low) to latch the 32 bits of data into the driver's latch. However, it isn't this latch that actually drives the segment.

To present an AC signal to the LCD, the backplane (BP) or common is driven by a clock signal. The segments need this clock as well.

The illusion of AC is accomplished by driving the segments with an XNOR of the BP clock and the latched data. The segment data tracks the clock whenever the latched data is 0. When the latched data is 1, the segment data is an inverted clock.

Segments being driven by an in-phase clock have no potential across them and appear white (LCD is transparent to light). Segments being driven by an out-of-phase clock have a reversing (AC) potential across them and appear black (LCD absorbs light).

The second routine selects which data is displayed. It starts by reading the RTC, which is one of the two options available on the PicStic (i.e., RTC or 12-bit ADC). A call to the RTC updates the year, month, day, day of week, hour, minute, and second variables.

I chose to use of all the registers by providing a rolling display of the day of the week (two letters), the month and day (MM.DD), the year (YYYY), and the hour and minute (HH:MM).

Each group displays for 1 s with the hour and minute hanging around for -10 s. The clock is continuously read until the seconds change, and then a decision is made about which group to display.

Here's an example of how the hour and minute are displayed. The minute byte variable (value 45h) holds two

BCD digits--4 and 5. These digits 4 and 5 are placed into temporary registers set aside for the third and fourth digit data.

The hour byte is handled the same way and stored into the first- and second-digit data registers. If the hour is greater than 12, 12 is subtracted because the first digit cannot be a 2.

Starting with the fourth digit, they are passed one at a time through the look-up table where the BCD digit value is replaced by a value corresponding to the segments enabled to display that BCD digit. Prior to sending the first-digit data (which is the last byte sent), the colon bit is masked so the time is displayed as HH:MM.

The other groups are handled in similar fashion except for the day of the week. A bit of cheating is done to display day of the week on a seven-segment display. Therefore, this routine has its own look-up table. To be honest, a bit of imagination must be used to interpret these (see Figure 5).

If the RTC came already set for life, I'd be finished here. But, there's the little matter of being able to set or change any one of the clock/calendar functions. And, this code is by far the most lengthy.

Two inputs are added-SW1 and SW2. These two inputs make it possible to select a register to change and increment through all the possible data choices.

When SW1 is depressed, the normal display routine is interrupted and execution passes to the update routine. The first RTC register is displayed. (Since the display routines are modular, I can take advantage of them in this update routine.)

Pressing SW2 increments this register, checks to make sure the maximum legal number isn't exceeded, resets the value to the minimum if it was, and redisplay the new entry. Whenever SW1 is pressed, the next register is selected. Once all the registers have been updated (SW1), the data is saved to the RTC and we go back to displaying real-time data.

CURRENT CURRENT

Currently, this circuit is drawing under 2 mA using a 5-V V_{CC} . You can

see that LCDs don't need the heavy currents required by LEDs.

This advantage is due to the fact that LCDs do not manufacture light—they only reflect or absorb it. Needless to say, they aren't useful at night or in a darkened area without an auxiliary source of light.

Electroluminescent panels have been a source of LCD backlight for years now, most recently as Timex's Indiglo. LED backlighting is becoming popular, but there goes the low current consumption.

LCDs do have their niches. They're in more and more consumables. I keep hearing that the hang-it-on-the-wall TV is only a few years away. But, the closest we've come are color LCDs in more expensive notebook computers.

Always happy to hear your views. If there's interest, I'll discuss the use of multiplexed displays in a future column. Till next time.. ☐

Special thanks to <<http://scorpio.kent.edu/~phil/cpip/lc/lc2.html>> and <<http://bly.colorado.edu/lc/lcintro.html>>.

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.

SOURCES

Varitronix VI-302-DP-RC-S, Microchip AY0438

Digi-Key Corp.
701 Brooks Ave. S
Thief River Falls, MN 56701-0677
(218) 681-6674
Fax: (218) 681-3380

PicStic 2
Micromint, Inc.
4 Park St.
Vernon, CT 06066
(860) 871-6170
Fax: (860) 872-2204

I R S

419 Very Useful
420 Moderately Useful
421 Not Useful



Starting at
\$139

Micromint's BlackJack-552 is a true "telecontroller" incorporating both microcomputer and FCC-certified Xecom modem in a single package. BlackJack comes preprogrammed with firmware utilities which are optimized for assembly language, C, and BASIC programs. Like Domino, BlackJack is easy to use. Attach power and a terminal, then upload, store, and execute your program.

SPECIAL FEATURES

- 80C552 processor with firmware monitor, 14.7456 MHz
- Serial I/O; (up to 38,400 bps); full-duplex TTL, I²C bus
- Connection: 1.2" dual-row 48-pin DIP header
- Parallel I/O: 20 bits, 4 bits shared with RTC
- Two interrupts: three timers, watchdog timer, two PWM outputs
- Hardware real-time clock
- A/D converter; 8 channels, 10 bits, 20k samples/sec.
- 2400 bps data modem (data/fax available)
- Power: +5V @ 55 mA;
Size: 2.75" x 1.375" x 0.5" potted
- 32K bytes SRAM, 32K bytes ROM and 128K bytes EEPROM
- Industrial temperature available



4 Park Street • Vernon, CT 06066


Call 1-800-635-3355

(860) 871-6170 • Fax (860) 872-2204

Haute Chips

SILICON UPDATE

Tom Cantrell



The fall fashion season is in full swing, and the recent Hot Chips VIII conference provided an opportunity for the hippest designers to strut their stuff.

Process geometries are down and clock rates are up. Ten-million-transistor chips definitely call for fabrics that breathe. Bell-bottom buses are back, and smart dressers accessorize with color-coordinated caches.

I don't mean to blaspheme, but it does seem there's a somewhat humorous similarity between chip and fashion design. Trends, whether they're hemlines or instruction sets, seem to come and go in rather whimsical cycles.

What was hot last season may be a "not" this year. But, you can't clean out your closet—it may be hot again next time.

From what little I've seen of fashion shows, some of the outfits are rather extreme. If I actually saw someone wearing something like that in an office or at the grocery store, I'm not sure if I'd laugh, faint, or call a cop.

Likewise, many of the chips at the Hot Chips show seemingly border on

the outrageous. Indeed, given the real-world bias of *INK* [and presumably, you, the readers], I predict almost none of you will design these chips in.

However, like at a fashion show, the goal is to discern the mainstream trends hidden beneath glitter, chains, and plastic wrap. Look closely, and you may find the runway features of today that will trickle down onto tomorrow's Main St.

BIG & TALL CPUs


Having seen firsthand the progression from monumental mainframes to desktop supercomputers, I've followed the triumphs and travails of computer architecture with some interest. Up until now, microprocessor designers have enjoyed the luxury of upgrading and refining historical computer-design tricks, taking advantage of increasing density and faster clock rates.

Lately, however, the follow-the-leader approach seems to be running out of gas. Microprocessors have incorporated, expanded, tweaked, and tuned just about every known concept to death.

The fact is, it seems to take an awful lot of transistors to speed things up a little. Fortunately, transistors are still getting cheaper and faster.

Whether that trend can continue is one big question. At the previous Hot Chips VII, no less an expert than Gordon Moore said that he has quit worrying about a wall, having mistakenly predicted its appearance many times before. At least for now, the operative assumption is silicon will march on. Transistors are fast and free, so go to it.


It's doing something useful with all those transistors that's the challenge. Here's where things don't look so rosy,



After visiting this year's Hot Chips show,

Tom has plenty to say about the increasing availability and speed of transistors. He discusses the latest tradeoffs between clock rate and complexity.

Figure 1—The good news is that faster clocks and more complex architectures are yielding better performance. The bad news is that progress on both fronts is increasingly difficult.



with obstacles on a number of fronts. Let's take a look at some of the whizziest chips, and you'll see what I mean.

CLOCK VS. COMPLEXITY

One challenge is finessing the increasingly inexorable tradeoff between clock rate and complexity. Figure 1, taken from an HP presentation, clearly illustrates the tradeoff. The Depth of Instruction Reorder window serves as an apt surrogate for complexity, while efficiency (SPECInt95/MHz x 1000) reflects CPI (clocks per instruction).

The HP PA-8000 and DEC 21164 illustrate the situation perfectly. The complex HP chip gets about twice as much done each cycle, but the cycles are about twice as long. So, when all's said and done, the HP and DEC chips end up delivering similar performance.

The problem addressed by reordering is simple enough. The good news is the latest superscalar designs are capable of executing up to four instructions every clock cycle. The bad news is they still execute zero instructions per clock when stalled.

Thus, the penalty for stall-inducing nasties like cache misses—already bad due to the growing gap between CPU and DRAM speeds—increases with the order of superscalarity.

Consider the typical-one instruction loads a register and the next uses the loaded value. If the load misses, a

200-MHz four-way superchip has to suck data from a 20-MHz DRAM, so the miss penalty is 10 clocks.

Thus, 40 instruction slots go whizzing idly by between the load and use instructions. The CPU only delivers a fraction of its theoretical performance and CPI more befitting an 805 1.

The goal of reordering is to ensure a pool of instructions is available at all times to keep the CPU fed. Consider the HP scheme in Figure 2.

Up to four instructions are fetched in order per cycle. These are divided into ALU and MEM IRBs (Instruction Reorder Buffers) that hold up to 28 instructions each.

From these pools, the CPU can execute up to four instructions per cycle out of order. Results are temporarily stored in so-called rename registers, which are hidden copies of the programmer-visible registers.

Though executed out of order, instructions (up to two each ALU and MEM) are retired (i.e., the programmer-visible registers are updated from the hidden rename registers) in order, allowing for precise (i.e., recovery is possible) exceptions.

Sounds simple enough, but there's a big gotcha in the form of dependencies. The CPU can't just execute instructions willy-nilly. The load-and-use scenario is one example of an operand dependency. Needless to say, it's bad

form to execute the use before the load!

Another example is PSW bits like carry. The CPU must not execute an instruction that depends on the carry bit before executing all prior instructions that update it.

Unfortunately, there is a long list of possible dependencies. The HP presentation has a list of half a dozen that ends with the somewhat plaintive sounding "Many others. . ." Fact is, keeping track of all the dependencies is horribly messy, which is why the reordering feature ends up consuming 850k transistors and a whopping 20% of the die area!

CACHE FLOW CRISIS

By contrast, the equally muscular DEC 21164 eschews reordering and overly esoteric architectural doodads in favor of blistering 400-MHz speed or faster. Rumor has it only those with a ham radio license can get a datasheet!

Not surprisingly, this approach isn't hassle free either. Those of you who deal with workaday micros may just hang a crystal on it and never give clocking another thought, but it's a big deal for RF-class chips.

It's not easy to ship a clock that goes 400+ MHz, while keeping skew in check. It calls for some big drivers, careful wiring, and a lot of power. Consider that the 21164 clock genera-

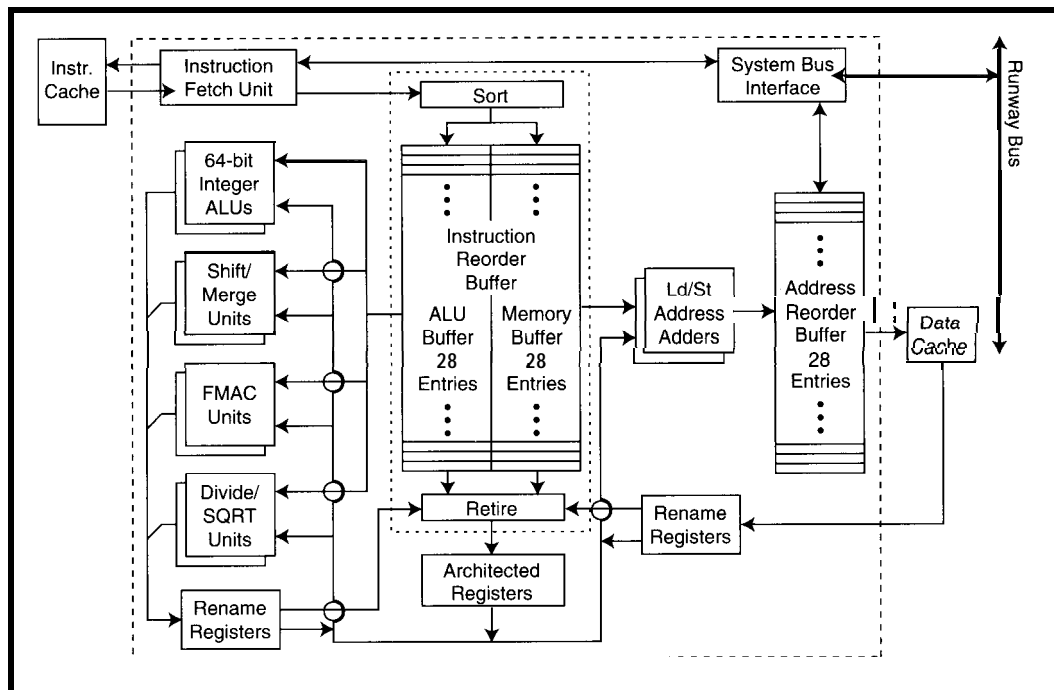
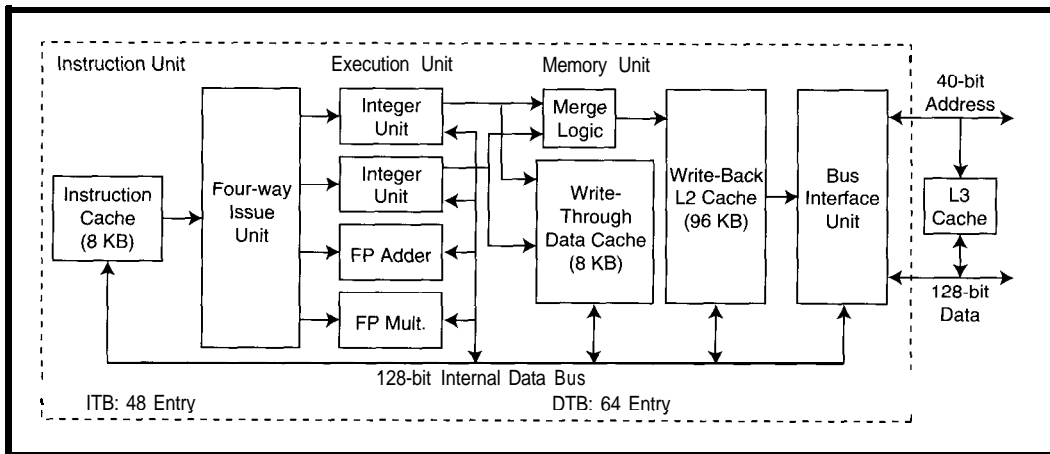


Figure 2-Though out-of-order execution isn't a new concept, now designers of chips like the HP PA-8000 have the luxury of applying brute force (i.e., lots of transistors) to the problem.

Figure 3—Running at 400+ MHz, the DEC Alpha 21164 relies on three levels of cache to fry to bridge the gap with slow DRAMs.



tor alone consumes a remarkable 15 W (-40% of the chip power)-and that's with the core running at only 2.5 V (3.3 V external).

Of course, boosting the clock rate into the stratosphere only makes the memory bottleneck worse, unless you happen to stumble across a source for 1-ns DRAMs.

The answer, naturally, is cache and plenty of it. As shown in Figure 3, the 21164 includes primary (L1) instruc-

	Power Reduction (W)
VDD Reduction	4.4X to 5.9
Reduce Functions	3x to 2.0
Scale Process	2x to 1.0
Clock Load	1.5x to 0.6
Clock Rate	1.25x to 0.5

Table 1—Applying their Alpha expertise to the simple ARM architecture, DEC comes up with a chip (StrongARM) that's both fast (160 MHz) and low power (<1 W).

tion and data caches (8 KB each) and a 96-KB secondary (L2) cache on the same die. Indeed, the 21164 even supports an optional L3 cache externally to more completely isolate the CPU from pokey DRAMs.

DEC also applies turbocharging techniques to their variant of the ARM RISC they call StrongARM. Some may say boosting the formerly mild-mannered ARM to 200+ MHz makes about as much sense as dropping a Corvette motor in a VW. But, how many skeptics would ante up their pink slips at the drag strip?

Remarkably, the chip consumes only half a watt [caveat: at $V_{CC} = 1.65-2.0$ V, see Table 1]. It even has 20-mW idle and 50- μ A sleep modes. The latter

cuts all power to the core and shuts off the PLL responsible for boosting a 3.68-MHz input clock to triple digits.

Yes, the chip has 2.5 million transistors, but most are devoted to cache. In fact, the CPU core itself (lacking bloaty superfeatures) only consumes a few hundred thousand transistors, which is less than the HP reorder buffers by themselves.

There's no doubt StrongARM wins the MIPS-per-watt war. The only question is whether there's much demand for battery-powered mainframes.

If clocks and complexity are hassles, why not dispense with both! The Amulet project from the University of Manchester (UK) purports to do just that. Like StrongARM, the complexity issue is dealt with by relying on a simple ARM CPU core. However, Amulet doesn't require a fast clock or, in fact, any clock at all!

The purpose of the ongoing project is to see whether asynchronous (no clock) design techniques make sense. In such a design, data doesn't flow through the chip in (c)lockstep. Instead, it relies on each functional unit to signal the need or ability to deliver data.

In theory, asynchronous design offers intriguing potential, especially for reducing power. Of course, there's no need for a hot-rod clock generator, but savings go beyond the obvious.

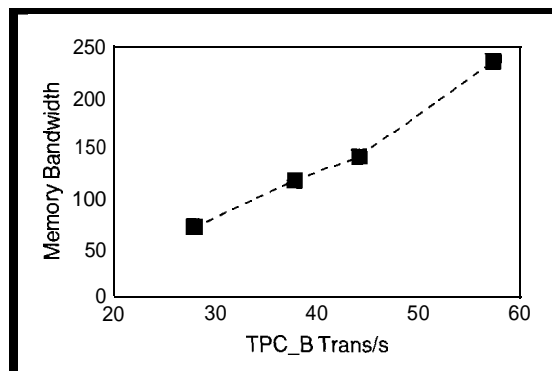
For instance, the increasingly popular technique of turning off the clock to unused functions is intrinsic to async design. Only units doing work consume power.

Not burdened with the concept of cycle time, the scheme can possibly improve performance as well. For instance, a multiply unit terminates early for easy calculations, and all transfers automatically exploit bursting characteristics (e.g., sequential cache lines hit faster)

Furthermore, async facilitates true cut-and-paste chip design. Independent function blocks can be combined at will, with no concern for clock skew. The clock solution (i.e., none) is the same whether combining 10, 100, or 1000 async units-something that certainly can't be said for synchronous designs.

Tired of the gadget on your bench hashing out your TV? Unlike the powerful continuous interference of a clock, async EM1 should be more broadband.

Figure 4—Too much focus on fancy CPU features may be misguided since there's some evidence that performance is primarily a function of main-memory bandwidth.



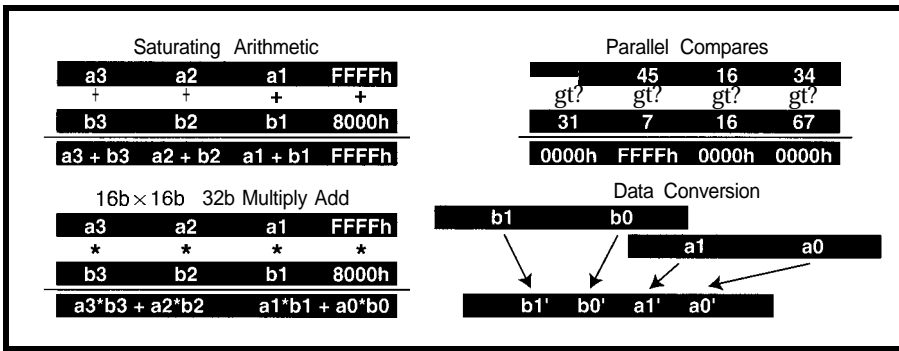


Figure C-Complex instructions are coming back, but this time, they crunch A/V data rather than the ASCII strings and BCD data of earlier CISCs.

However, there are a few tricky problems with the concept, most revolving around the lack of determinism. In principle, the design may be correct, but proving it (via design verification) is quite difficult.

Production test and debug become interesting challenges—just where will you connect that scope trigger? Finally, there's the small matter of pitching nonreal-time computers that run at different (and variable!) speeds.

The team has managed to fabricate a few chips, jolly good show that. They do report that "...traces show some

unpredictable behavior," but that malady is not uncommon. At this point, it appears keeping up with sync performance is the challenge, but the jury is still out on power reduction.

RAM CRAM

Let's turn our attention back to that old nemesis—main-memory bandwidth.

A group at UC Berkeley (including RISC luminary, Dave Patterson) appears to essentially wave the white flag in terms of architectures' continued ability to finesse around the in-


creasing gap between fast CPUs and slow DRAMs. Throwing ever more cache at the problem has been the easy way out, but it can't overcome skyrocketing (100s of clocks) miss penalties.

One of my own pet peeves is that the superchips with giant caches that work so well in marketing benchmarks seem to crumble in the face of real applications. The presentation cited an article in which a state-of-the-art CPU delivered a meager 12% of theoretical bandwidth on a database application.

They also showed evidence that instruction sets, pipelining, superfeatures, cache, compilers, and so forth are all well and good, but bottom-line system performance is largely predicted by one factor. You guessed it—main-memory bandwidth (see Figure 4).

The UCB proposal to bridge the CPU and DRAM gap is real simple. Just stick 'em both on the same chip!

In fact, examples of just such chips (called IRAMs for Intelligent RAMs) are starting to emerge from the re-



PLIX:

- * performs all X-10 transmit and receive functions
- * detects AC power loss
- * connects to parallel port
- * provides simple interface timing
- * uses very little power (only 1.8 mA @ 5 V)

WE'VE SOLD THOUSANDS!

TRY IT OUT! TO GO ALONG WITH THE PLIX CHIP AND ITS DATA SHEET, WE HAVE:

- * an easy-to-use self-powered PC parallel port interface kit
- * sample control software

FIRE UP A PLIX AND BLAST X-10 OVERHEAD OFF YOUR SYSTEM!

PLIX Evaluation kit (includes PLIX chip and evaluation board) — \$39
 PLIX chip ————— \$20 (single qty.), \$12 (100 qty.)
 TW523 X-10 powerline transceiver ————— **\$30**
 (Prices do not include shipping)

CALL 1-800-635-3355 TO ORDER!




MICROMINT, INC.

4 Park Street • Vernon, CT 06066

Tel: (860) 871-6170 • Fax: (860) 872-2204 • <http://www.micromint.com/>
 in Europe: (44) 1285-658122 • in Canada: (514) 336-9426

Distributor Inquiries Welcome!

EASY TO USE CAD



Now in Both Windows and DOS

CAD PAK Windows-\$199, DOS - \$159

Ideal for New Users, Hobbyists, & Small Businesses!
 Provides everything for PCB Layout & Schematic Drawing

PROTEUS Starts at \$425

Most Powerful EDS System on the market!!!

- * Schematic Capture * Circuit Simulation
 - * PCB Layout with Rip-Up & Retry Autorouting
- ALL PRICES ARE IN US FUNDS. PLB INCLUDE \$109/HR

R4 SYSTEMS INC.

1100 GORHAM ST. Suite 11B-332
 NEWMARKET ONTARIO
 CANADA L3Y 7V1

905 898-0665 FAX 905 898-0683
 BBS 905 898-0508 (9600,8,N,1)

FREE DEMO
 WRITE OR CALL
TODAY

Download Demo - <http://www.r4systems.on.ca>
Internet email info@r4systems.on.ca

search labs. Consider a Mitsubishi chip presented at the last ISSCC conference that combines a RISC CPU with a whopping 8 MB of DRAM (see Photo 1).

Testimony to the concept is Motorola's recent agreement to license Mitsubishi's DRAM-plus-logic process.

Putting DRAM on the same chip not only eliminates those messy (and slow) drivers, pins, traces, and so on, but it also allows the DRAM to be organized as small blocks for speediest access. Nevertheless, it's quite possible (as the Mitsubishi chip illustrates) to include a complement of cache on-chip as well.

Though conceptually appealing, there are quite a few real-world challenges to the scheme. Most basic is the fact that optimized DRAM and logic processes are somewhat like oil and water—they don't mix well.




Figure 6—With 2–4x performance boost, MMX may allow the CPU to take over the tasks formerly handled by dedicated modem, audio, and graphics chips.

Consider that the typical SRAM cell on a conventional single-chip micro (i.e., logic process) may be 10 or more times the size of that on a pure DRAM chip. Likely as not, silicon wizardry combined with the inexorable (we all hope) increase in density will chip away at the manufacturing concerns.

There are a few potential problems noted, such as the simple fact you can't add more memory to a particular IRAM after the fact. This may not be

that big a deal if the chip has a gigabit of DRAM. Even if that's not enough, the CPU will be such a small part of the die size (already well less than half for today's cache-heavy chips), you won't mind chucking it in the drawer with your obsolete SIMMs.

It may be quite likely that IRAMs will be a hit down the road. In the short run, however, I

think there's plenty of mileage to be extracted from more conventional solutions such as SDRAM.

Also, when presenters talk of a shift in the balance of power between the DRAM and micro markets, they tend to overlook just how different the two businesses are. The micro biz is a complex and technical sale, calling for a lot of evangelism, support, tools, and patience. By contrast, DRAMs are a price-and-delivery commodity, more akin to pork bellies than computers.

ENGINEERING PROFESSORS LISTEN UP!

The reason you read *Circuit Cellar INK* is obvious. Like thousands of others, you have come to realize that *Circuit Cellar INK* presents technically relevant computer-based applications in a comprehensive and instructive manner. In fact, our projects are presented just the way you'd do them if only you had the time.


Since 1994, we have provided thousands of college students with industry-quality development tools and information. This program is **FREE** and because of the continued success we have experienced, we'd like to extend our offer to the 1996-1997 school year.

If you're an engineering professor or teach qualified technical students, we'd like to give you and all of your students free subscriptions to *Circuit Cellar INK*. Please tell us how we can contact you directly so you don't miss a single issue. Just a reminder: **IF YOU WERE A PARTICIPATING COLLEGE LAST YEAR, PLEASE BE SURE TO UPDATE US WITH YOUR CURRENT ROSTER EACH SEMESTER.**

WE'RE HELPING AMERICA REGAIN THE COMPETITIVE EDGE. WE'RE HOPING YOU WILL TOO!

Rose Mansella, Circulation Coordinator
Circuit Cellar INK
4 Park St.
Vernon, CT 06066

Tel: (860) 875-2199 Fax: (860) 871-0411
Internet: rose.mansella@circellar.com
World Wide Web: <http://www.circellar.com/>



Lowest Cost Data Acquisition

ADAC's new Value-Line has uncompromising design features and high quality components at prices below the low cost guys!

Just check out the specs:

Lowest Cost

5500MF

8 channels 12-bit A/D,
16 digital I/O, Counter/Timer

\$195

High Speed

5508LC

8 channels 12-bit A/D,
1 00KHz, DMA

\$245

Multi-Function DMA

5516DMA

16 channels 12-bit A/D,
DMA, 16 digital I/O

\$295

High Resolution

5500HR

16 channels 16-bit A/D,
DMA, 8 digital I/O

\$595

learn more:

voice 800-648-6589
fax 617-938-6553
web www.adac.com
email info@adac.com

ADAC

American Data Acquisition Corporation
70 Tower Office Park, Woburn, MA 01801 USA




Photo 1--The M32R/D combines a 32-bit RISC, 2-KB cache, and 2 MB of DRAM. According to Mitsubishi, the internal DRAM offers nearly three times the speed, yet it consumes only about half the power of external DRAMs.

FORWARD TO THE PAST?

Yes, the "clock and power" problem, "memory bandwidth" problem, and (lack of available) "instruction-level parallelism" problem are converging into one big headache.

Less cynical readers may actually see irony in the latest weapon in the architect's arsenal—namely, new and more complex instructions.

Creeping CISCism has turned into a veritable flood. For example, additions on the previously mentioned RISCs include memory-prefetch and branch-prediction hints (PA-8000), multiply and accumulate (StrongARM), and extensions to directly support bytes and words (21164).

But, that's just the start. The big bloat in instruction sets is in the form of various multimedia extensions pioneered by HP, picked up by Sun (VIS), and poised to penetrate PCs in the form of Intel MMX.

MMX, the forthcoming multimedia extension to the 'x86 architecture, is like the others in concept. All rely on new instructions (e.g., MAC, saturating math, and format conversions) that operate in an SIMD/vector-like parallel fashion on packed A/V data (see Figure 5). Heck, them thar complex instructions ain't so bad after all, as you see in Figure 6.

Guess I better make room in the closet—I think there's a bit of space between the Nehru jacket and disco getup—for those VAX manuals, Forth disks, and C-machine dissertations.

You just never know. ☐

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 by E-mail at tom.cantrell@circellar.com, by telephone at (510) 657-0264, or by fax at (510) 657-5441.

SOURCES

Hot Chips

1658 Belvoir Dr.
Los Altos, CA 94024
(415) 941-6699
Fax: (415) 9415048
<http://www.hot.org/hotchips/>

PA-8000

Hewlett-Packard
Portable Computer Div.
1000 NE Circle Blvd.
Corvallis, OR 97330
(541) 757-2000

DEC 21164, StrongARM

Digital Semiconductor
(510) 490-4752
<http://www.digital.com.info/semiconductor/>

M32R/D

Mitsubishi Electronics
1050 E. Arques Ave.
Sunnyvale, CA 94086
(408) 730-5900
Fax: (408) 730-4972

422 Very Useful
423 Moderately Useful
424 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

From time to time, I like to review for new subscribers or recently connected readers the various services we have available via the Internet and the Circuit Cellar BBS. In addition to giving you a bit of our history our Web site provides subscription info, introduces our staff, and announces contests, among other things.

Through our BBS, we offer several computer engineering discussion groups and full E-mail access to the rest of the Internet. This and a few engineering-related news groups greatly enhance your ability to learn what the latest high-tech solutions are.

We look forward to soon being able to provide you with direct access to the BBS via the Internet. We'll keep you posted on these developments.

The following message threads may be obtained in their entirety either from the BBS or on disk. See "Article Software" at the end of this column for more information.

Pressure sensors

This thread starts with the question, "Does anybody know of a low-cost pressure sensor for measurement of barometric pressure?"

A quick search of an index by one of our other users results in a list of 132 company names. Needless to say, I can't run the complete thread here. However, the list is available on-line for anyone who wants to check it out.

The remainder of the thread mentions a few of the more promising companies by name as a place to start the search.

As part of the same thread, the original author also posts a follow-up question: "Do you know of a shuttle or motorized valve that can be used to switch air to a pneumatic on/off valve?" Again, a number of useful company names come up.

PID Control

The question that sparked this thread is, "Does anyone know of a source of generic or specific info about PID control programming?" Of course, the response lists many available references, from application-oriented primers to theory-heavy treatises. Even past **Circuit Cellar INK** articles are mentioned.

On realizing that a generic treatment of PID algorithms is far too complex for the task at hand, the original poster

details what he's trying to do. He has a device that receives fanfold paper from a printer and restacks it. The output from the printer is very jerky, so he wants his device to adapt to the rate at which the paper is coming from the printer.

The result is a lengthy discussion of various solutions to the problem.

File deletion

The final summarized thread this month starts off with, "During the process of learning a few things about assembly language, I have created invalid filenames. DOS and Windows can see the files, but I am unable to delete them. They say these invalid files don't exist. Is there a trick to wiping these files off my hard drive?"

The general consensus among BBS users includes several tricks. First, reformat the drive (a bit drastic). Next, try using wildcard characters when referring to the files in rename or delete commands. Finally, use a utility to modify the names or manually delete the files. Norton Utilities is especially suited to the task, but even DOS Debug can be used in a pinch.

Wire inductance

Msg#: 3865

From: Lyndon Walker To: All Users

Can anyone give me a ballpark estimate of the inductance and capacitance of four miles of 22-24-AWG copper pair, either twisted or not, probably hung from telephone poles? Signal is 9600-bps current loop.

I already called a wire manufacturer, but my experience has been that when a support person says, "I'll have to ask Engineering and call you back," I could be waiting for a while.

Msg#: 3906

From: Jan Verhoeven To: Lyndon Walker

The British mathematician Heavyside made this subject his lifework. Perhaps looking up his name (or that of Maxwell, his major opponent who was believed by everyone, although Heavyside was correct) in an encyclopedia could solve your quest?

CONNECT TIME

Msg#: 3909

From: Lyndon Walker To: Jan Verhoeven

There *must* be more recent stuff than this.

I tried calculating it, but I made a bunch of assumptions that are probably incorrect, so I don't trust the answers. I just need to get in the ballpark so I can build a model to test my idea.

I'm going to use the equations I found for calculating inductance of coax cable in an old physics text. I hope that's close enough so it has a chance of working the first time out in the field.

Msg#: 3973

From: Pellervo Kaskinen To: Lyndon Walker

I don't think you need to be concerned about inductance if you are using a line with relatively high resistance and a current loop. Let's see..

For 22 AWG, two wires four miles end-to-end is 680-Ω loop resistance. At 20-mA current, you need 13.6 V just to overcome the line losses plus something for your receiver. At 24 AWG, the resistance in the loop is 1080 Ω. Then the voltage drop is 21.6 V.

However, if you still want to check the inductance, then you should also check the capacitance.

In fact, we really are talking about a transmission line in the classical sense. In such a case, what matters is the characteristic impedance:

$$Z = \sqrt{\frac{L}{C}}$$

In this equation, L is the inductance per any fixed length and C is the capacitance per the same length.

Now, rather than calculating both the inductance and the capacitance separately, it is common to use a formula that produces the Z directly. In the two parallel wires far away from the ground (up in the poles), the applicable formula is:

$$Z = \frac{120}{\sqrt{k}} \times \ln \frac{2D}{d}$$

where:

k is relative dielectric constant for the material between the wires, 1 for air, about 2.8 to 3.3 for most plastics;
 \ln is the natural logarithm;

\$149.⁹⁵
QTY. 1

SINGLE BOARD COMPUTER



This board has the I/O and Processor power to handle demanding applications: Serial, A/D, D/A, LCD, Keypad, Real Time Clock, EEPROM, and Flash. This board has got it all at a price that is out of this world!

- * 8 High-Drive Outs & 16 Programmable Digital I/O Lines
- * 8 Channels of 10/ 12 Bit A/D & Optional 4 Channels of 12 bit D/A
- * 2, 16 bit Timers & 2 optional 16 bit Counters
- * Up to 3 RS232/485 Serial Ports
- * Backlit Capable LCD Interface
- * Optional 16 Key Keypad & Interface
- * 1Meg of Memory Space Total, 32K ROM & 32K RAM Included
- * **Assembler & Monitor Included, BASIC and Forth Optional**

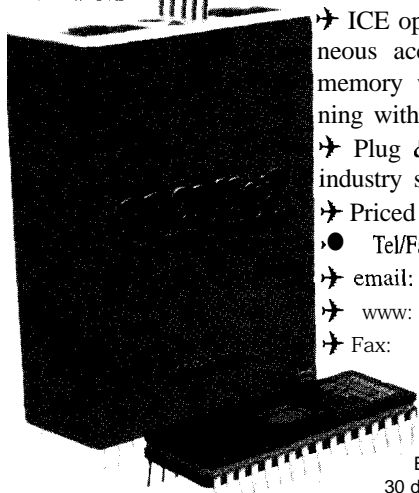
1985-1996
OVER
10
YEARS
OF SERVICE

EMAC, inc.

618-529-4525 Fax 457-0110 BBS 529-5708
P.O. BOX 2042, CARBONDALE, IL 62902

PROMJet[®]-ICE

- Ultra compact EPROM and FLASH emulator with highest download speed (1-4 Mb/S), largest memory capacity (1~32Mb) and fastest access time (85~25ns) in the industry
- Other features include 3V target support, jumperless configuration, battery backup, 128 bit bus support and external power supply. ➤ Fits directly into memory socket or use extension cable for flexibility. ➤ Compact design based on high density FPGAs and double-sided surface-mounted 4 layer PCB for added reliable operation
- ICE option allows simultaneous access to PROMJet memory while target is running without waitstate signals
- Plug & Play drivers for industry standard debuggers
- Priced from US\$295 / MB
- Tel/FaxBack: 206.337.0857
- email: emutec@emutec.com
- www: www.emutec.com
- Fax: 206.337.3283



2.2x1.7x0.7"

EmuTec Inc
 Everett Mutual Tower
 2707 Colby Av, Suite 900
 Everett, WA 98201, USA
 30 day money-back policy
 Visa & Mastercard accepted

CONNECT TIME

D is the distance between the wires;
 d is the wires' diameter **in the same units** as the distance.

As an example, the 22-AWG wires mounted 6" apart in open air, yields 737 Ω as the characteristic impedance. If they are like a cable with a plastic insulator, they are probably right next to each other.

With the D reducing and k increasing, it is no surprise that the characteristic impedance becomes something in the 115-130-R range for most commercial paired cables. With special low-capacitance formulations (foam insulation), you may get over 150 Ω , and with the twin lead used for the TV antenna connections, you get 240 or 300 Ω .

If you are worried about the need to "drive" an inductive circuit, no sweat. The capacitance in the cable compensates the inductance, and vice versa. You do not need high voltage to drive a fast-rising edge into the cable, nor do you see the capacitance eating the front edge of your signal.

What you do see is a traveling time called group delay. The lower the characteristic impedance, the slower the wavefront travels. We may be talking about something like 60% of the speed of light.

I hope this answers your questions. If not, give more details of your tentative choices and somebody here is bound to do the "sanity check" for you.

Like, with a four-mile stretch of wires on poles, what are you doing to protect the systems at both ends against over-voltages that any thunderstorm is going to cause? Plus, are you planning a unidirectional or a bidirectional link?

Msg#: 3984


From: Lyndon Walker To: Pellervo Kaskinen

This is just what I was wondering about. I had **no** idea how much inductance to expect, so I didn't want to just shrug it off.

I didn't even consider the line capacitance until going through the **Radio Amateur's Handbook** hoping to find an inductance table. This all came about because when testing the idea, I messed up and the output waveform from the optocoupler was rather rounded. This triggered the thought that I might have to worry about things other than line resistance.

As far as lightning goes, I'm thinking of the Teccor Sidactors I used in a previous modem project.


Low Cost CAD Software
 for the IBM PC and Compatibles
 Now In WindowsTM95



- Easy to use schematic entry program (**SuperCAD**) for circuit diagrams, only \$149. Includes netlisting, bill of materials, extensive parts libraries.
- Powerful, event-driven digital simulator (**SuperSIM**) allows you to check logic circuitry quickly before actually wiring it up. Works directly within the **SuperCAD** editor from a **pull-down** menu and displays results in "logic analyzer" display window. Starting at \$149 this is the lowest cost simulator on the market. Library parts include TTL, and CMOS devices.
- Analog simulator (**mentalSPICE**) for \$149. Allows AC, DC and transient circuit analysis. Includes models of transistors, **discretes**, and op amps.
- Circuit board artwork editor and autorouter programs (**SuperPCB**), starting at \$149. Produce high quality artwork directly on dot matrix or laser printers. You can do boards up to 16 layers including surface mount. Includes Gerber and Excellon file output. Autorouter accepts **netlists** and placement data directly from the **SuperCAD** schematic editor.
- Low cost combination packages with schematics and PCB design: P-layer for \$399, 16-layer for \$649.

Write or call for free demo disks:
MENTAL AUTOMATION, INC.
 5415 - 136th Place S.E.
 Bellevue, WA 98006
 (206) 641-2141 • BBS (206) 641-2846
 Internet: <http://www.mentala.com>


386 SBC \$83



- OEM (1K) PRICE
 INCLUDES:
 • 5 SER (8250 USART)
 • 3 PAR (32 BITS MAX)
 • 32K RAM, EXP 64M
 • STANDARD PC BUS
 • LCD, KBD PORT
 • BATT. BACK. RTC
 • IRQ0-15 (8259 X2)
 • 0237 DMA 8253 TMR
 • BUILT-IN LED DISP.
 • UP TO 8 MEG ROM
 • CMOS NVRAM
 USE TURBO C,
 BASIC, **MASM**
 RUNS DOS AND
 WINDOWS
 EVAL KIT \$295

\$95 SINGLE
 PIECE
 PRICE
**UNIVERSAL
 PROGRAMMER**

- DOES 8 MEG EPROMS
- CMOS, EE, FLASH, NVRAM
- EASIER TO USE THAN MOST
- POWERFUL SCRIPT ABILITY
- MICROCONT. ADAPTERS
- PLCC, MINI-DIP ADAPTERS
- SUPER FAST ALGORITHMS



OTHER PRODUCTS:
 8088 SINGLE BOARD COMPUTER OEM \$27 ... • 95
 PC FLASH/ROM DISKS 128K-16M) 21 75
 18 BIT 18 CHAN ADC-DaC CARD 55 195
 WATCHDOG (REBOOTS PC ON HANGUP) 27 95

- EVAL KITS INCLUDE MANUAL BRACKET AND SOFTWARE.
- 5 YR LIMITED WARRANTY
- FREE SHIPPING
- HRS: MON-FRI 10AM-6PM EST

MVS MVS BOX850
 MERRIMACK, NH
 (508) 792 9507

CONNECT TIME

Ultrasonic sensors

Msg#: 3571

From: Larry Siu To: All Users

I would like to locate the source of an ultrasonic sensor which has a frequency response of 20-50 kHz or sensors with center frequencies of 25 kHz and 50 kHz \pm 5kHz. I need these sensors to detect the presence of 25 kHz and 50 kHz in the environment. If anyone knows where to get these parts, please drop me a note. Any pointers will be appreciated.

Msg#: 3581

From: Ken Simmons To: Larry Siu

What about the ubiquitous pinger used in Polaroid cameras?

Tom Cantrell just finished a wonderful article on the "PID Pong" toy a couple of months ago (see *INK* 42, 50, 69, and 70), and it uses that kind of ultrasonic sensor.

Msg#: 3634

From: Chris Boros To: Larry Siu

Just a note about the Polaroid transducers. They do handle frequencies from about 20 kHz to 100 kHz (at least the old instrument grade unit did), but are *highly* directional. That should not be surprising, considering their intended use in range finding.

If you want a nice wide-angle transducer, try one of the Panasonic units [I know they make a 25-kHz unit, not sure about 50 kHz], which are available from Digi-Key.

Msg#: 3748

From: Jan Verhoeven To: Chris Boros

"Wide angle" and "ultrasonics" are mutually exclusive features! The shorter the wavelength of the sound, the lower the tendency to go along curved trajectories.

So, if you make a "nice wide-angled ultrasonic transducer," you're gonna be punished by the power requirements.

Msg#: 3818

From: Larry Siu To: Steve Ciarcia

I guess I need to explain what I want to do. I need to pick up a rat's chirps in a specific frequency, namely 25 kHz and 50 kHz. I plan to either built bandpass filters to single out the frequency above or build a frequency-to-voltage converter.

Either way, I need to know the sensitivity of the transducer and the output voltage. Can the Polaroid unit do the job?

Msg#: 3819

From: Steve Ciarcia To: Larry Siu

A quick call to Polaroid will get you a datasheet on their transducer. Of course, only you know how much noise the rat actually makes to know if it will be sensitive enough.

Msg#: 3923

From: Pellervo Kaskinen To: Jan Verhoeven

> "Wide angle" and "Ultrasonics" are mutually exclusive

Agree? Not quite! Let's look at some of the issues a little closer.

In 'principle * , what makes a narrow angle is the planar shape of the radiating surface in most transducers. Or, you might say that any "point source" is wide angle to the extent of being omnipolar.

On any larger surface, the wave behavior according to Huygens produces a narrow beam due to the interaction of all the elementary waves from different spots on the radiating surface. The elementary waves subtract from each other in directions where the distances are not equal. They add in the one direction where the distances are equal.

Then, the sad *practical* point. If we try to make wider-angle radiation available by reducing the dimension of the transducer, we need more power to overcome the reducing coupling efficiency from the solid surface to the air.

This, of course, can only be taken so far. When the transducer size is reduced toward a point source, the energy density approaches infinity. Clearly an unsupportable situation. Is there any remedy?

A couple of suggestions to consider are using baffles right in front of the transducer and making the transducer surface spherical rather than planar. I know that both approaches are possible to some extent. The baffles, cones, or horns actually are used in some equipment I have seen advertised in the past.

I have never seen the spherical shape implemented. I don't think there are enough people wanting the wide angle for the manufacturers to become interested. And of course, you cannot make it a full sphere for several reasons that I don't care to even start pondering.

But, there is a reasonably practical way of living with these physical constraints. If a single transducer does not provide a wide enough pattern, use an array!

Mount a dozen or so small transducers closely next to each other on a portion of a spherical surface. Connect them in the same phase, parallel or series, and feed all of them from the same source. You get a fair approximation of a wide-angle pattern on an ultrasonic transducer.

CONNECT TIME

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 <<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.

Message threads summarized at the beginning of the column are also available on the BBS for at least six months after they are first posted. The subject line at the start of each summary matches the subject used on the messages themselves. Simply call the BBS and

search for those subject lines to find the message threads.


For those unable to download files or messages, the software and messages are also available on disk. Software for issues prior to 1995 comes on a 360-KB IBM PC-format disk, one issue per disk. For issues from 1995 on, software comes on a 1.44-MB PC-format disk, with three issues per disk. Disks cost just \$12 each. 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) 875-2199. Be sure to specify the issue numbers with your order. Please add \$3 for shipping outside the U.S.

425 Very Useful

426 Moderately Useful

427 Not Useful

PIC it μ P!



The PIC-MDS is a complete training and development environment for the mid-range PIC™ microcontrollers. Proven in classrooms, the PIC-MDS lets you learn to use and

develop applications for PIC microcontrollers quickly and easily. Our step-by-step examples guide you from simple programs through to keypad scanning, serial I/O, LCD display, ND conversion, data logging, and interrupt routines.

The PIC-MDS Professional Pack includes:

- PIC16C71 & PIC16C84 microcontrollers
- EPIC Programmer & PM assembler
- detailed training manual with disk
- in-circuit programming cable & socket
- ZIF socket • buffered LED port indicators
- 2 X 16 LCD display • 2 analog pots • RS-232 port
- 256 byte serial EEPROM • crystal/resonator socket
- 5V & variable DC power supply • AC voltage adapter
- all I/O pins on screw terminals & PICBUS connector
- all code examples in Microchip and Parallax™ syntax

only \$299!

Also Available:

- Professional Board (no programmer) **\$259**
 - Hobbyist Pack **\$199** (kit, no ZIF, '84 only)
 - Hobbyist Board **\$139** (kit, no ZIF, '84 only, no programmer)
 - Teachers manual **Call**
- Prices exclude shipping. Parallax is a registered trademark of Parallax, the PIC & Microchip are trademarks of Microchip Technology Inc.

Sirius microSystems

Visit our Web page: <http://www.siriusmicro.com>

172 Harvard Road
Waterloo, ON N2J 3V3
CANADA
+1 519 886 4462 (ph.)
+1 519 886 4253 (fax)
support@siriusmicro.com

Parts List Manager

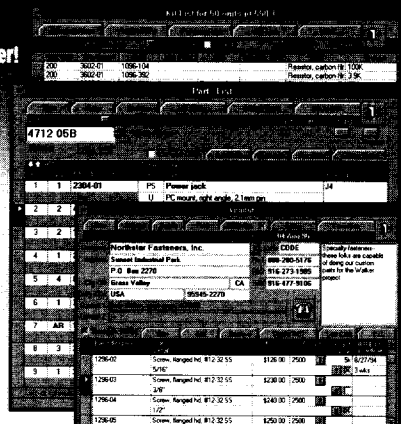
For Engineers, Product Designers, Prototypers

Windows-based software helps you stay organized ...And makes the job easier!

and manage multi-assembly parts lists for products in development ...And after!

Keep track of:

- Part Specs & Dwg's
- Suppliers & Mfrs
- Boms and Kit Lists
- Product Costs
- Engineering Stock



Parts Vendors™

Version SE: \$99 + s/h Extended: \$299 + s/h Rev. Hist: \$399 + s/h

Call 800-280-5176



Trilogy DESIGN

Requires 486, 10 meg min. ram. Win 3.x or Win95

916-273-1985
fax 916-477-9106
P.O. Box 2270, Grass Valley, CA 95945
<http://www.trilogymicro.com>

CONNECT TIME

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.

360-KB

1.44-MB5 on, software comes on a PC-


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 <<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

IRS

425 y Useful

PIC it μ P!



PIC™ microcontrollers. Proven in classrooms, the PIC-MDS lets you learn to use and develop applications for PIC microcontrollers quickly and easily. Our step-by-step examples guide you from simple programs through to keypad scanning, serial I/O, LCD display, A/D conversion, data logging, and interrupt routines.

The PIC-MDS Professional Pack includes:

- PIC16C71 & PIC16C84 microcontrollers
- EPIC Programmer & PM assembler
- detailed training manual with disk
- in-circuit programming cable & socket
- ZIF socket • buffered LED port indicators
- 2 analog pots • port
- crystal/resonator socket
- 5V & variable DC power supply • AC voltage adapter
- all I/O pins on screw terminals & PICBUS connector
- all code examples in Microchip and Parallax™ syntax

only \$299!

Also Available:

- Professional Board (no programmer) **\$259**
 - Hobbyist Pack **\$199** (kit, no ZIF, '84 only)
 - Hobbyist Board **\$139** (no programmer)
 - Teachers manual **Call**
- Prices exclude shipping Parallax is a registered trademark of Parallax, Inc. PIC & Microchip are trademarks of Microchip Technology Inc.

Sirius microSystems

Visit our Web page: <http://www.siriusmicro.com>

172 Harvard Road
Waterloo, ON N2J 3V3
CANADA
+1 519 886 4462 (ph.)
+1 519 886 4253 (fax)
support@siriusmicro.com

Parts List Manager

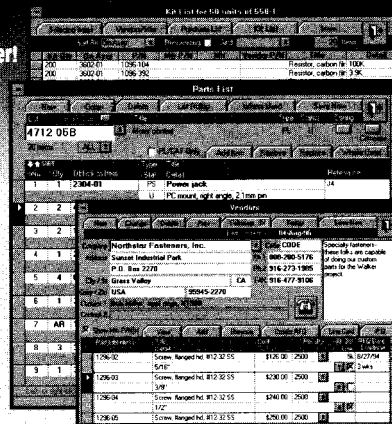
For Engineers, Product Designers, Prototypers

Windows-based software helps you stay organized ...And makes the job easier!

Easily create and manage multi-assembly parts lists for products in development ...And after!

Keep track of:

- Part Specs & Dwg's
- Suppliers & Mfrs
- Boms and Kit Lists
- Product Costs
- Engineering Stock



Parts & Vendors™

NEW! VERSION 2.0

Version SE: \$99 + s/h Extended: \$299 + s/h Rev. Hist: \$399 + s/h

Call 800-280-5176



Trilogy DESIGN

916-273-1985

fax 916-477-9106


P.O. Box 2270, Grass Valley, CA 95945

Requires 486, 10 meg min. ram. Win 3.x or Win95

<http://www.trilogydesign.com>

PRIORITY INTERRUPT

Hardly a High-Definition Opinion



I'm sure you've been in a situation where a subject has come up and your first reply is, "Don't get me started on that!"

Whenever I say that, it's usually a topic I feel very strongly about and that most assuredly has a polarizing effect on others within earshot. For most people, the "verboten" subjects are things

Bill Clinton.

While I might have equally acerbic opinions concerning these topics, the subjects that give me a ballistic disposition these days are Microsoft, HDTV, FCC, IRS, and pretty much anything else the Feds have their noses in.

Personal experience usually elevates a subject on to or off of my list. "Software" used to be one of the words that always elicited a response. I've since conceded that it's a necessary component of embedded control. My concern these days is with companies who routinely fleece the development community with their latest, barely tested, flimflam upgrade.

The FCC and IRS have probably been on my list the longest. Every time I cut the FCC some slack, they pull out a new regulation that either guarantees a monopoly for a small but significant bunch of companies or makes an outlaw out of every electronics experimenter who ever twisted two wires together. The IRS, on the other hand, has no regulation favorites. They simply presume everyone's an indentured servant.

The real paradox arises when a bunch of these entities begin to squabble and I'm forced, heaven help me, to feel sympathetic toward one of them. The issue that evokes this mercy is the new HDTV standard.

Basically, it's a fight between entertainment and computer manufacturers. The entertainment industry has finally decided that 50 years of **525-line** NTSC analog television is enough, and they've proposed a **1080-line** digital high-definition TV standard. The transmission technique still involves sending an interlaced scanned image, but with four times the bandwidth. Presumably, it retains some compatibility with present TVs, so they all don't become instant junk.

The computer industry advocates progressive "noninterlaced" scanning on a lower-resolution image. The idea is to make it easier to scan, compress, and manipulate video into packets for use in computers. Somewhere in the middle of the argument are the cinematographers who want more resolution but, at the same time, are becoming increasingly dependent on computer-generated video production. They're pushing not only for HDTV itself but wider aspect ratios than the presently proposed **16:9** (NTSC is **4:3**).

If it were purely economic, the decision might be easier. Computer software comes from the U.S. and TV sets are made elsewhere, right? Lets sell computers and software? Well, not quite. Virtually all new **HDTVs** will be made here (albeit by European and Asian companies). OK, let's create new jobs and sell the entertainment industry HDTV.

This is but a small snippet of the argument. Any resolution will have its down side. The FCC has a real problem with this one, and I'm not sure there's any one best solution. All I know is that one industry is 50 years old and the other is barely 15.

The advances we've seen in the computers video compression, communication speed, and processing power has allowed it to accommodate today's video standard. Why shouldn't continued technological advances accommodate the next standard as easily? Handicapping the entertainment and cinematography business simply to fit the present view of computer technology impugns the impulse that drives the whole computer industry.

The alternative is for Microsoft to announce a standard and sell an endless series of HDTV beta-level protocols to broadcasters and manufacturers (at exorbitant prices), while the rest of us throw out our TVs. At the same time, we'd get an opportunity to upgrade or replace our PCs for the new Windows HDTV (at exorbitant prices) and subscribe to the Microsoft cable/satellite/telephone/ether communication service (at exorbitant prices). And, of course, it will be news according to MSNBC. Who else?

steve.ciarcia@circellar.com