PROJECT: Audio Sensor Software
LOCATION: Italy
PAGE: 14

INSIGHT: Conformité Européenne Design Conformity
LOCATION: France
PAGE: 46

ANALYSIS: Measure Stepper Motor Performance
LOCATION: United States
PAGE: 58

CIRCUIT CELLAR

THE WORLD'S SOURCE FOR EMBEDDED ELECTRONICS ENGINEERING INFORMATION

DECEMBER 2011
ISSUE 257

PROGRAMMABLE LOGIC

MCU-Based Sleep-Stage Analysis

DIY Smart Electronic Load Design

Embedded Linux System Platforms

Interference Immunity for Electronic Designs

Joystick Control with Bluetooth Connectivity

PLUS

Micro Design
Inside the Mind of a Microprocessor Developer

// Analog Bipolar Tech to Digital Design
// The Usefulness of Verilog HDL
// Design Using Programmable Logic
// And More



7 25274 75349 9
$7.50 U.S. ($8.50 Canada)
www.circuitcellar.com

The easiest finished product you will ever create

Start with the PK70 EX



Design your own, or use a NetBurner Personality Blade



Add your label to get a finished product



Personality Blades

Prototype Blade *(P/N NBPKBP-100CR)* ..... **$25 ea.**
Create your own Personality Blade

CPLD Digital I/O Blade *(P/N NBPKD-100CR)* ..... **$59 ea.**
Programmable Xilinx CPLD
32 digital I/O

Multi-I/O Blade *(P/N NBPKM-100CR)* ..... **$99 ea.**
Eight 12-bit analog-to-digital inputs
Two 16-bit digital-to-analog outputs
16 digital I/O

FPGA Digital I/O Blade *(P/N NBPKX500-100CR)* ..... **$169 ea.**
Programmable Xilinx Spartan 3E FPGA
High-density 62-pin connector

NTP Server with GPS Blade *(P/N NBPKX500-100CR)* ..... **$99 ea.**
Network Time Protocol (NTP) server
Precision GPS time reference

Quad UART Blade: RS-232 Version
(P/N NBPKBU-232CR) ..... **$99 ea.**
Network enable up to 4 RS-232 serial devices

Quad UART Blade: RS-422/485 Version
(P/N NBPKBU-485CR) ..... **$119 ea.**
Network enable up to 4 RS-422/485 serial devices

Quad UART Blade: RS-232/422/485 Version
(P/N NBPKBU-MMSCR) ..... **$245 ea.**
Network enable up to 4 RS-232/422/485 serial devices

PK70 EX & Development Kit

PK70 EX *(P/N PK70EX-100IR)* ..... **$199 ea.**
32-Bit Freescale ColdFire 147 MHz CPU
8MB SDRAM • 4MB Flash Memory
SD/MMC Flash Card Support • I2C • SPI • Address and Data Bus
10/100 Ethernet • 2 UARTs

PK70 EX Development Kit
(P/N NNDK-PK70EX-KIT) ..... **$299 ea.**
The Development Kit features NetBurner's Eclipse, an enterprise level professional IDE that provides editing, downloading and debugging in one environment.

NetBurner
Networking in One Day!

Information and Sales | sales@netburner.com **Web** | www.netburner.com **Telephone** | 1-800-695-6828

NEW chipKIT™ Development Platforms from Microchip

The first 32-bit MCU platforms compatible with Arduino™ HW and SW



The chipKIT™ Uno32 and Max32 development boards are the first 32-bit microcontroller-based platforms that are compatible with existing Arduino™ hardware and software.

The chipKIT platform allows hobbyists and academics from many disciplines, such as mechanical engineering, computer science and artists to develop original embedded applications easily and quickly including: motor control, LCD display, wired and wireless communications, LED matrix control and sensor networks

Key Features:

- Application development using an environment based on the original Arduino IDE, modified to support PIC32 devices while still supporting the original Arduino line. Leverages existing code examples, tutorials and resources.
- Pin-out compatibility with many existing Arduino shields
- Higher performance at a lower price-point than existing solutions
- Advanced capabilities including:
  - Integrated USB (Device/Host, OTG)
  - Integrated Ethernet
  - CAN

Development Board Comparison

Feature	Core	Performance	Program Memory (KB)	RAM (KB)	Additional Features
chipKIT™ Uno32	32-bit	80 MHz	128	16	PMP/PSP/RTCC
chipKIT Max32	32-bit	80 MHz	512	128	USB, 2x CAN, Ethernet, DMA, RTCC

IT'S EASY TO GET STARTED!

- Visit **www.microchip.com/chipkit**
- Purchase a chipKIT development board
- Download the free software
- Order free PIC32 samples



chipKIT Uno32 (TDGL002)



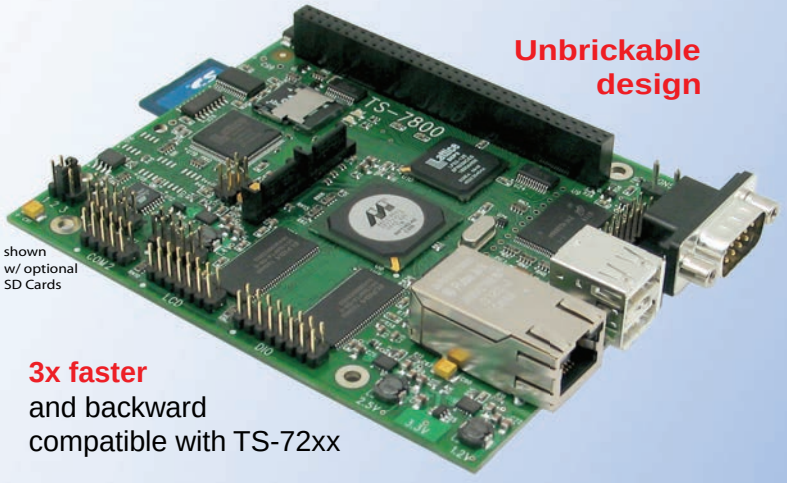
chipKIT Max32 (TDGL003)

The Microchip name and logo, the Microchip logo and PIC are registered trademarks and chipKIT is a trademark of Microchip Technology Incorporated in the USA and other countries. All other trademarks mentioned herein are the property of their respective trademark holders. These trademark holders are not affiliated with Microchip Technology Incorporated, and do not support, sponsor or endorse chipKIT products or solutions. © 2011 Microchip Technology Incorporated. All rights reserved. 5/11

MICROCHIP

Embedded Systems

High-End Performance with Embedded Ruggedness



shown w/ optional SD Cards

Unbrickable design

3x faster and backward compatible with TS-72xx

TS-7800
500 MHz ARM9

- Low power - 4W@5V — $229 qty 100
- 128MB DDR RAM
- 512MB high-speed (17MB/sec) onboard Flash — $269 qty 1
- 12K LUT customizable FPGA
- Internal PCI Bus, PC/104 connector
- 2 host USB 2.0 480 Mbps
- Gigabit ethernet
- 2 SD sockets
- 10 serial ports
- 110 GPIO
- 5 ADC (10-bit)
- 2 SATA ports
- Sleep mode uses 200 microamps
- Boots Linux 2.6 in 0.7 seconds
- Linux 2.6 and Debian by default

TS-TPC-7390
7" Color Touch Panel Computer

- Low Power, Industrial Quality Design
- Mountable aluminum frame
- 200 MHz ARM9
- 64MB SDRAM (128MB opt)
- 512MB Flash w/ Debian Linux
- 800 x 480 video core
- Dedicated framebuffer- 8MB RAM
- Audio codec with speaker
- Boots Linux 2.6 in less than 2 seconds
- Unbrickable, boots from SD or NAND
- Runs X Windows GUI applications
- Runs Eclipse IDE out-of-the-box



$449 qty 1

More Touch Panel Computers on our website

- Over 25 years in business
- Open Source Vision
- Never discontinued a product
- Engineers on Tech Support
- Custom configurations and designs w/ excellent pricing and turn-around time
- Most products stocked and available for next day shipping

Design your solution with one of our engineers (480) 837-5200

New Products

TS-WIFIBOX-2

A Complete Solution for 802.11g WiFi Applications

NEW!

$185 qty 1



- Low power (3.2 watts), fanless
- Power via 5-12 VDC, USB, PoE (opt.)
- 64MB DDR-RAM
- 256MB ultra-reliable XNAND drive
- Micro-SD Card slot
- RS-232, RS-485, CAN, RTC
- Header with SPI and 11 DIO
- 480Mbit/s USB, Ethernet, PoE option
- Boots Linux 2.6.24 in < 3 seconds
- Un-brickable, boots from SD or flash
- Customizable FPGA - 5K LUT
- Optional DIN mountable enclosure

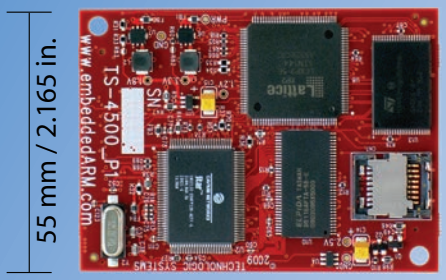
Ideal for gateway or firewall, protocol converter, web server, WiFi audio, and unattended remote applications

TS-SOCKET Macrocontrollers

Jump Start Your Embedded System Design

NEW!

series starts at $92 qty 100



- Dual 100-pin connectors
- Secure connection w/ mounting holes
- Common pin-out interface
- Low profile w/ 6mm spacing

TS-SOCKET Macrocontrollers are CPU core modules that securely connect to a baseboard using the TS-SOCKET connector standard. COTS baseboards are available or design a baseboard for a custom solution with drastically reduced design time and complexity. Start your embedded system around a TS-SOCKET Macrocontroller to reduce your overall project risk and accelerate time to market.

Current projects include:

- TS-4200: Atmel ARM9 with super low power
- TS-4300: Cavium ARM11 with dual 600 MHz and FPU
- TS-4500: Cavium ARM9 at very low cost
- TS-4700: Marvell PXA168 with video and 1.2 GHz CPU
- TS-4800: Freescale iMX515 with video and 800 MHz CPU
- Several COTS baseboards for evaluation & development

Technologic Systems

We use our stuff.

Visit our TS-7800 powered website at

www.embeddedARM.com

TASK MANAGER

Embedded Community

Up for a quick game? Let's play with the words **circuit** and **cellar**:

> If you enjoyed reading about **The Circuit Cellar** in Steve Ciarcia's **"Circuit Cellar"** column back in the 1980s, then today you likely read **Circuit Cellar, Inc.**'s ***Circuit Cellar*** magazine in your own **circuit cellar**.

Are those typos? Let me explain.

What do you think of when you see the word **circuit** next to the word **cellar**? What's the difference between **The Circuit Cellar** and **a circuit cellar**? What if you capitalize the first letter of each word and put them in quotes (i.e., **"Circuit Cellar"**)? What do you think of when you remove the quotes (i.e., **Circuit Cellar**)? What if you italicize the two words (i.e., ***Circuit Cellar***)?

OK, enough of the editor's game. I have an actual point to make. During the last three decades, the term **circuit cellar** has referenced many different, yet related, things. And, yes, punctuation, grammar, and context all matter.

First, note that **The Circuit Cellar** has always been, and still is, a physical location at Steve Ciarcia's house. I've been there. It's full of electronics, old parts, and cool gadgets.

Next, there's **"Circuit Cellar"** with quotation marks. In the mid-'80s, the term referred to Steve's famous column in *Byte* magazine. Later, after Steve left *Byte*, the term quickly took on a few new meanings all at once. When italicized, the term referred to the internationally respected magazine you're reading today: ***Circuit Cellar***. And an entity had to publish the magazine, right? Hence the company name **Circuit Cellar, Inc.**

And then there's **circuit cellar**. Once Circuit Cellar, Inc. started reaching thousands of readers around the globe with its content, engineers began referring to their own workspaces as circuit cellars. That's right. Years before the term "hackspace" became fashionable, designers were using what they'd read in *Circuit Cellar* to complete innovative projects in their own circuit cellars.

As we wrap up another great year of publishing, I encourage you to reflect on the various ways the term **circuit cellar** has evolved during the past 30 years. It has come to mean so many important things to so many people, from professionals to academics. This positive idea—the notion that our company supports and educates the embedded design community—motivates our staff to strengthen our content and broaden our offerings each year.

On behalf of the staff, I thank all the authors who contributed articles during the past 12 months. I also thank our members who read *Circuit Cellar* each month, chat with our editors, and interact with our authors. I appreciate your interest and support.

FYI: I purposely referred to you as our **members** rather than **subscribers**. Why? You are truly part of a community—The Circuit Cellar Community—and it's a fellowship you help make fresh and exciting every day.

The Circuit Cellar Community? Hmm. I think we now have another term to discuss in 2012. Let me know what you think.

cj@circuitcellar.com

C. Abate

CIRCUIT CELLAR®

THE WORLD'S SOURCE FOR EMBEDDED ELECTRONICS ENGINEERING INFORMATION

FOUNDER/EDITORIAL DIRECTOR
Steve Ciarcia

EDITOR-IN-CHIEF
C. J. Abate

ASSOCIATE EDITOR
Nan Price

CONTRIBUTORS
Jeff Bachiochi
Bob Japenga
Robert Lacoste
George Martin
Ed Nisley
Richard Wotiz

PROJECT EDITORS
Ken Davidson
David Tweed

PUBLISHER
Hugo Van haecke

ASSOCIATE PUBLISHER
Shannon Barraclough

CUSTOMER SERVICE
Debbie Lavoie

CONTROLLER
Jeff Yanco

ADMINISTRATIVE COORDINATOR
Valerie Luster

ART DIRECTOR
KC Prescott

GRAPHIC DESIGNER
Grace Chen

Circuit Cellar's mission is to collect, select, and disseminate need-to-know information around the world in the fields of embedded hardware, embedded software, and computer applications. Circuit Cellar uses an assortment of print and electronic content-delivery platforms to reach a diverse international readership of professionals, academics, and electronics specialists who work with embedded, MCU-related technologies on a regular basis. Our aim is to help each reader become a well-rounded, multidisciplinary practitioner who can confidently bring innovative, cutting-edge engineering ideas to bear on any number of relevant tasks, problems, and technologies.

ADVERTISING
800.454.3741 • 978.281.7708 • www.circuitcellar.com/advertise

ADVERTISING REPRESENTATIVE
Peter Wostrel
Strategic Media Marketing, Inc.
1187 Washington St., Gloucester, MA 01930 USA
800.454.3741 • 978.281.7708
peter@smmarketing.us • www.smmarketing.us
Fax: 978.281.7706

ADVERTISING COORDINATOR
Kim Hopkins
E-mail: khopkins@circuitcellar.com

Cover photography by Chris Rakoczy—Rakoczy Photography
www.rakoczyphoto.com

PRINTED IN THE UNITED STATES

CONTACTS

SUBSCRIPTIONS
Information: www.cc-access.com, E-mail: subscribe@circuitcellar.com
Subscribe: 800.269.6301, www.cc-access.com, Circuit Cellar Subscriptions, P.O. Box 5650, Hanover, NH 03755-5650
Address Changes/Problems: E-mail: subscribe@circuitcellar.com

GENERAL INFORMATION
860.875.2199, Fax: 860.871.0411, E-mail: info@circuitcellar.com
Editorial Office: Editor, Circuit Cellar, 4 Park St., Vernon, CT 06066, E-mail: editor@circuitcellar.com
New Products: New Products, Circuit Cellar, 4 Park St., Vernon, CT 06066, E-mail: newproducts@circuitcellar.com

AUTHORIZED REPRINTS INFORMATION
860.875.2199, E-mail: reprints@circuitcellar.com

AUTHORS
Authors' e-mail addresses (when available) are included at the end of each article.

CIRCUIT CELLAR® (ISSN 1528-0608) is published monthly by Circuit Cellar Incorporated, 4 Park Street, Vernon, CT 06066. Periodical rates paid at Vernon, CT and additional offices. **One-year (12 issues) subscription rate USA and possessions $45, Canada/Mexico $60, all other countries $63. Two-year (24 issues) subscription rate USA and possessions $80, Canada/Mexico $110, all other countries $116.** All subscription orders payable in U.S. funds only via Visa, MasterCard, international postal money order, or check drawn on U.S. bank. **Direct subscription orders and subscription-related questions to Circuit Cellar Subscriptions, P.O. Box 5650, Hanover, NH 03755-5650 or call 800.269.6301.**

Postmaster: Send address changes to Circuit Cellar, Circulation Dept., P.O. Box 5650, Hanover, NH 03755-5650.

Circuit Cellar® 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® disclaims any responsibility for the safe and proper function of reader-assembled projects based upon or from plans, descriptions, or information published by Circuit Cellar®.

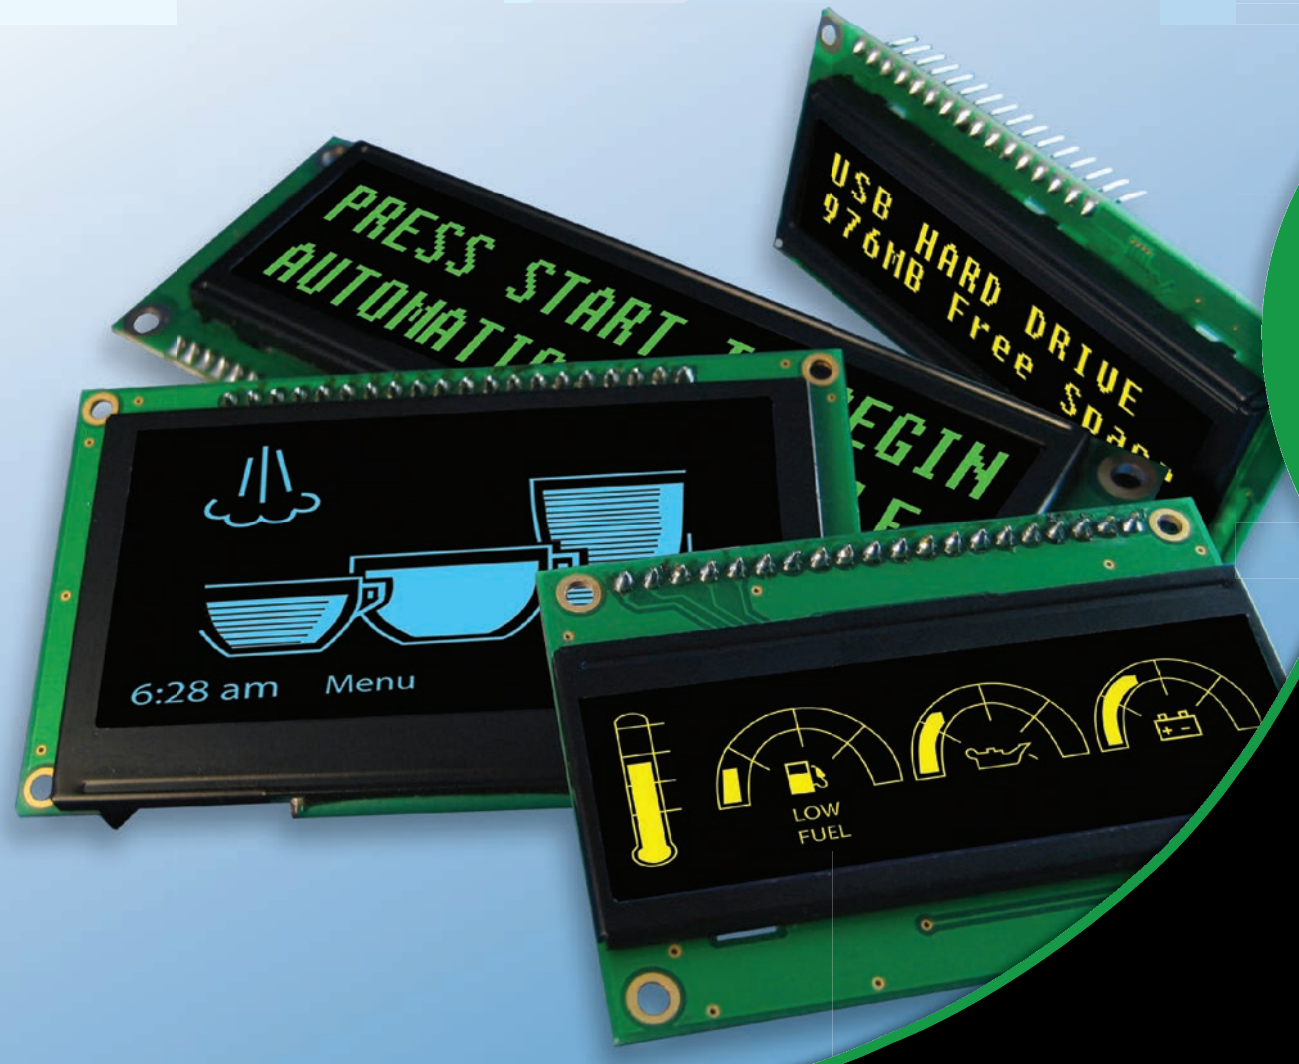
The information provided by Circuit Cellar® is for educational purposes. Circuit Cellar® makes no claims or warrants that readers have a right to build things based upon these ideas under patent or other relevant intellectual property law in their jurisdiction, or that readers have a right to construct or operate any of the devices described herein under the relevant patent or other intellectual property law of the reader's jurisdiction. The reader assumes any risk of infringement liability for constructing or operating such devices.

Entire contents copyright © 2011 by Circuit Cellar, Incorporated. All rights reserved. Circuit Cellar 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.

NEWHAVEN DISPLAY
INTERNATIONAL

New
Character and Graphic OLED Displays

- Eco-Friendly Design
- Self-Illuminated
- Faster response
- Better contrast
- Wider viewing angles



NEWHAVEN DISPLAY INTERNATIONAL

Newhaven Display's OLED (Organic Light Emitting Diode) is a brighter, higher contrast display that has faster response times, wider viewing angles and consumes less power than conventional LCD, LED or VFD Displays. OLED displays are self-illuminating so no backlight is required for maximum visibility in all environments. This also allows OLEDs to be significantly thinner than standard VFD, LED or LCD displays. Newhaven Display offers Character and Graphic type OLED displays as complete easy-to-use modules.

For more information visit www.newhavendisplay.com or contact us at 847-844-8795

RoHS Compliant

INSIDE ISSUE 257

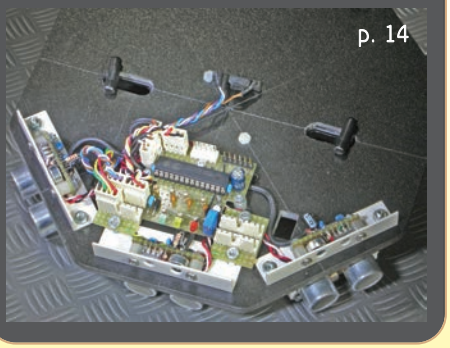
December 2011 • Programmable Logic

14 **Sound Tone Detection with a PSoC (Part 2)**
The Software & Sensing
Guido Ottaviani

24 **The (VI)sualizer**
A Smart Electronic Load
Curt Terwilliger

38 **Sleep-Stage Alarm Clock**
An MCU-Based System for Processing EEG Data
Michael Bohn, Tristan Money, Richard Neden, & Keita Sasaki

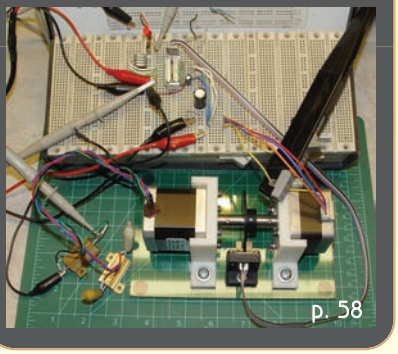
Sound Board & Sensors for Programming



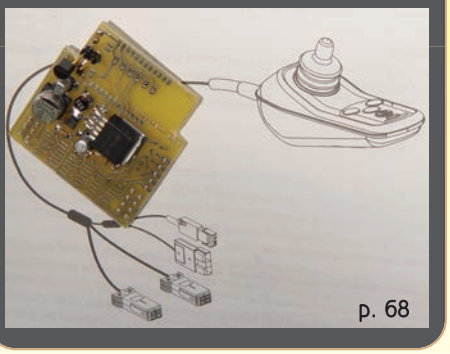
DIY Smart Electronic Load



Dynamometer with Two Steppers



Control System Design



46 **THE DARKER SIDE**
CE Marking
A Process to Ensure Product Conformity
Robert Lacoste

54 **EMBEDDED IN THIN SLICES**
Getting Started with Embedded Linux (Part 2)
Choosing a Platform for Your System
Bob Japenga

58 **ABOVE THE GROUND PLANE**
Stepper Motor Drive (Part 3)
Torques
Ed Nisley

64 **THE CONSUMMATE ENGINEER**
EMI and Data Integrity
George Novacek

68 **FROM THE BENCH**
Fly-By-Wire Wheelchair (Part 2)
Bluetooth Wheelchair Control
Jeff Bachiochi

TASK MANAGER 4
Embedded Community
C. J. Abate

TEST YOUR EQ SOLUTIONS 13

NEW PRODUCT NEWS 8

QUESTIONS & ANSWERS 20
IC Design Engineering and Beyond
An Interview with Monte Dalrymple

CROSSWORD 74

INDEX OF ADVERTISERS 79
January Preview

PRIORITY INTERRUPT 80
40 Years and Still Going Strong
Steve Ciarcia

AS9120A
Authorized Distributor

mouser.com

Find It Here. Faster.™

The Newest Products for Your Newest Designs®

Scan Here
to Watch Video

Authorized distributor for the most advanced
semiconductors and electronic components.
Get What's Next. Right now at mouser.com.

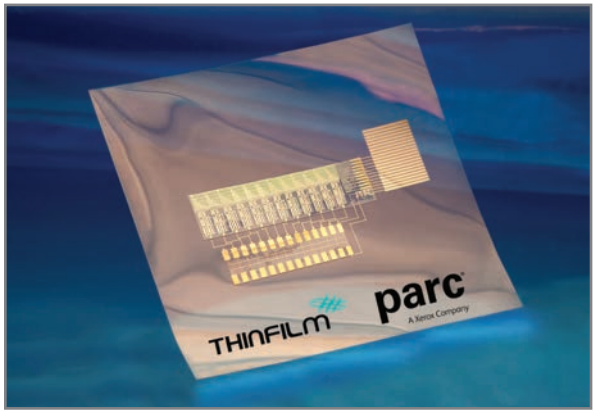
MOUSER ELECTRONICS®
a tti company

Mouser and Mouser Electronics are registered trademarks of Mouser Electronics, Inc.

SCALABLE PRINTED CMOS MEMORY

Thinfilm Addressable Memory is a printed nonvolatile memory device that combines Thinfilm's polymer-based memory technology with transistor technology from Palo Alto Research Center (PARC) using complementary pairs of N-type and P-type transistors to construct the circuits. The addition of the integrated circuits makes the roll-to-roll printed memory devices addressable by printable logic.

Thinfilm's Addressable Memory helps keep things connected via smart tags. Smart tags require the commercial availability of devices that are affordable, have rewritable memory, support integration with sensors and other electronic components, are environmentally friendly, and can be produced using high volume, roll-to-roll printing. The Thinfilm Addressable Memory prototype, rewritable memory with logic circuitry, will meet all of these requirements.



One target market for Thinfilm Addressable Memory products is near-field communications (NFC) tags, which are available in Android phones. NFC tags enable device-to-device communication. Other target markets include sensor tags and disposable price labels. Thinfilm Addressable Memory can be integrated with other printed components, such as antennas and sensors, to create fully printed systems for interaction with everyday objects. The memory devices can monitor the temperature of food and drugs or track retail items individually rather than by pallet, container, or truckload with a simple tap of an NFC-enabled phone.

Using printing to manufacture electronics minimizes the number of process steps, which in turn reduces manufacturing costs and lowers the environmental impact compared with traditional semiconductor processes. Thinfilm devices are thinner than traditional silicon devices and can be produced in form factors as slim as a strand of hair.

Contact Thinfilm for pricing.

Thin Film Electronics ASA (Thinfilm)
www.thinfilm.se

GRAPHICAL JTAG/BOUNDARY SCAN PROJECT DEVELOPMENT

CASCON Mission Assist is a graphical user interface (GUI) for SYSTEM CASCON, which is GOEPEL's JTAG/Boundary Scan software platform. CASCON Mission Assist enables completely graphical project development based on predefined mission targets, while automating the context-sensitive adaptation of all system tools and process controls. Intuitive system controls help improve a first-time user's productivity by safely guiding them through the project development flow. The open and flexible design of CASCON Mission Assist provides support for various access technologies, such as Boundary Scan, processor emulation test, chip-embedded instruments, in-system programming, and core-assisted programming based on a consistent software platform. This multidimensional instrumentation also features optimized test coverage and throughput provided by a definition of mixed testing and programming strategies.



The Mission Assist visualizes the entire project development workflow based on a selected mission target. This includes automated process scripting and complete management of centralized project data.

The new graphical project development based on Mission Assist is available with SYSTEM CASCON version 4.6.

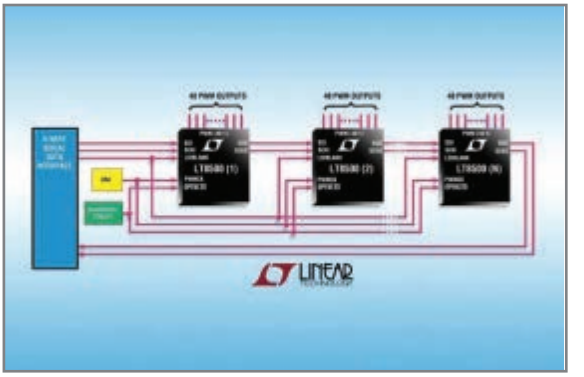
Contact GOEPEL for pricing.

GOEPEL electronics, Ltd.
www.goepel.com

NEW PRODUCT NEWS

PWM GENERATOR CONTROLS LEDs OR PROVIDES 48 DAC OUTPUTS

The **LT8500** is a pulse-width modulation (PWM) generator with 48 independent channels. Each channel has an individually adjustable 12-bit (4,096 step) PWM register and a 6-bit (64 step) ±50% correction register. All controls are programmable via a simple TTL/CMOs 50-MHz serial data interface. The LT8500 can be used for a variety of PWM-intensive purposes, such as light-emitting diode (LED) control, as well as industrial and robotic applications. For example, the LT8500 can be used with three LT3595As, 16-channel LED drivers, to deliver 48 independently dimmed LED strings for locally dimmed LED backlighting for large liquid crystal displays (LCDs). The correction resister enables matching light output of each LED string.



The LT8500 operates from a 3-V to a 5.5-V input range, enabling it to operate from a range of standard power buses. For LED applications, each channel has individual 6-bit dot correction current adjustment and 12-bit grayscale PWM dimming. The 40-ns on-time of the LT8500 offers wide dynamic contrast ratios. Both correction and grayscale dimming are accessible via a serial interface in TTL/CMOS logic.

The LT8500 offers two diagnostic information flags: synchronization error and open LED. The flags are sent with additional state information on the serial data interface during status read back. The 50-MHz cascadable serial data interface includes buffering and skew balancing, easing the design for PWM-intensive applications such as large-screen LCD dynamic backlighting. The combination of minimal externals and a 5 mm x 9 mm quad flat no-lead (QFN) package provides a compact-footprint solution for multichannel PWM applications.

Pricing for the LT8500EUHH, which is available in a 56-lead 5 mm x 9 mm QFN package, starts at **$2.95** each in 1,000-piece quantities. The extended temperature version, the LT8500IUHH (or I-grade), is also available, with pricing starting at **$3.25** each in 1,000-piece quantities.

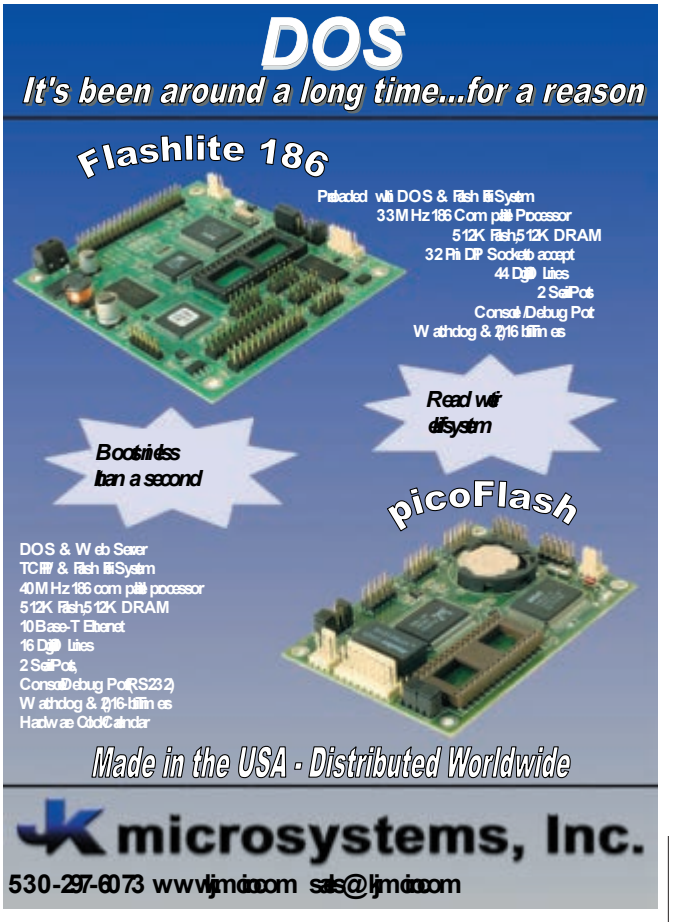
Linear Technology Corp.
www.linear.com

NPN



*ezlcd-301

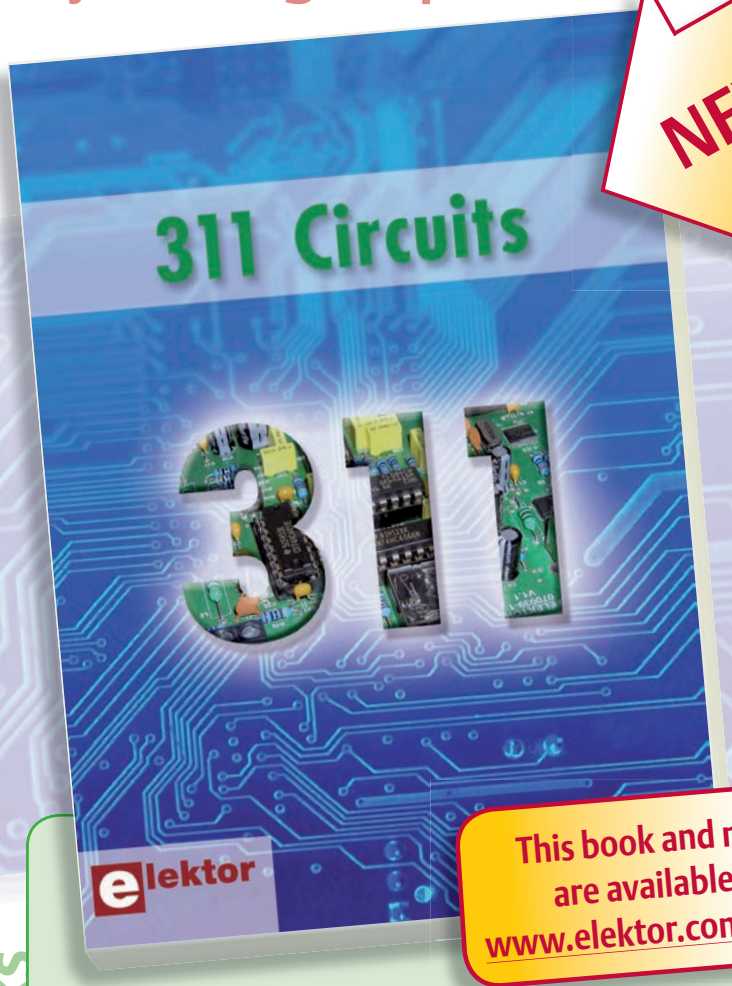
EARTHLCD.COM
We Make LCDs Work™
For more information visit http://store.earthlcd.com/cc1011



BOARDS, BOOKS, DVDs AND MORE AT WWW.ELEKTOR.COM/SHOP

Elektor Shop

The world of electronics at your fingertips!



NEW!

This book and more are available at www.elektor.com/books

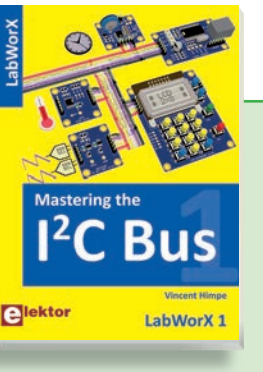
Books

Creative solutions for all areas of electronics

311 Circuits

311 Circuits is the twelfth volume in Elektor's renowned 30x series. These Summer Circuits compilation books have been bestsellers for many years. This brand new book contains circuits, design ideas, tips and tricks from all areas of electronics: audio & video, computers & microcontrollers, radio, hobby & modelling, home & garden, power supplies & batteries, test & measurement, software, not forgetting a section 'miscellaneous' for everything that doesn't fit in one of the other categories. 311 Circuits offers many complete solutions as well as useful starting points for your own projects. Both categories and anything in between represent a veritable fountain of inspiration for cultivating your own ideas and learning about electronics. This book deserves a place not far from the workbench!

420 pages • ISBN 978-1-907920-08-0 • $47.60

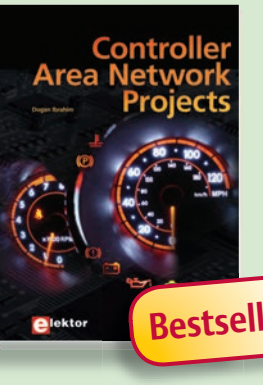


LabWorX: Straight from the Lab to your Brain

Mastering the I²C Bus

Mastering the I²C Bus is the first book in the LabWorX collection. It takes you on an exploratory journey of the I²C Bus and its applications. Besides the Bus protocol plenty of attention is given to the practical applications and designing a solid system. The most common I²C compatible chip classes are covered in detail. Two experimentation boards are available that allow for rapid prototype development. These are completed by a USB to I²C probe and a software framework to control I²C devices from your computer.

248 pages • ISBN 978-0-905705-98-9 • $47.60



Bestseller!

Free mikroC compiler CD-ROM included

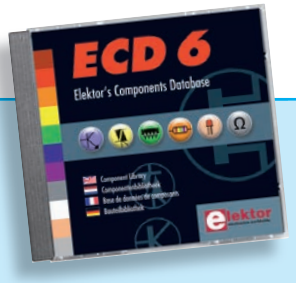
Controller Area Network Projects

The aim of the book is to teach you the basic principles of CAN networks and in addition the development of microcontroller based projects using the CAN bus. You will learn how to design microcontroller based CAN bus nodes, build a CAN bus, develop high-level programs, and then exchange data in real-time over the bus. You will also learn how to build microcontroller hardware and interface it to LEDs, LCDs, and A/D converters. This book comes complete with a free CD-ROM.

260 pages • ISBN 978-1-907920-04-2 • $47.60

Prices and item descriptions subject to change. E. & O.E

CD/DVD-ROMs



Elektor's Components Database

CD ECD 6

This CD-ROM gives you easy access to design data for over 70,000 components! The program package consists of eight databanks covering ICs, transistors, diodes and optocouplers. A further eleven applications cover the calculation of, for example, zener diode series resistors, voltage regulators, voltage dividers and AMV's. A colour band decoder is included for determining resistor and inductor values. All databank applications are fully interactive, allowing the user to add, edit and complete component data. This CD-ROM is a must-have for all Electronics enthusiasts!

ISBN 978-90-5381-258-7 • $40.20

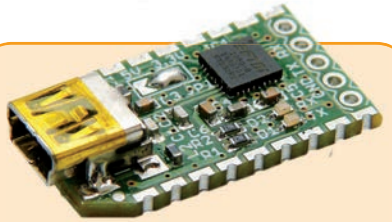


Circuits, ideas, tips and tricks from Elektor

CD 1001 Circuits

This CD-ROM contains more than 1000 circuits, ideas, tips and tricks from the Summer Circuits issues 2001-2010 of Elektor, supplemented with various other small projects, including all circuit diagrams, descriptions, component lists and full-sized layouts. The articles are grouped alphabetically in nine different sections: audio & video, computer & microcontroller, hobby & modelling, home & garden, high frequency, power supply, robotics, test & measurement and of course a section miscellaneous for everything that didn't fit in one of the other sections. Texts and component lists may be searched with the search function of Adobe Reader.

ISBN 978-1-907920-06-6 • $55.70



FT232R USB/ Serial Bridge/BOB

You'll be surprised first and foremost by the size of this USB/serial converter — no larger than the moulded plug on a USB cable! And you're also bound to appreciate that fact that it's practical, quick to implement, reusable, and multi-platform — and yet for all that, not too expensive! Maybe you don't think much of the various commercially-available FT232R-based modules. Too expensive, too bulky, badly designed, ... That's why this project got designed in the form of a breakout board (BOB).

PCB, assembled and tested

Art.# 110553-91 • $20.90

Elektor is more than just your favorite electronics magazine. It's your one-stop shop for Elektor Books, CDs, DVDs, Kits & Modules and much more!
www.elektor.com/shop

Elektor US
PO Box 180
Vernon, CT 06066
USA
Phone: 860-875-2199
Fax: 860-871-0411
E-mail: order@elektor.com



USB Long-Term Weather Logger

This stand-alone data logger displays pressure, temperature and humidity readings generated by I²C bus sensors on an LCD panel, and can run for six to eight weeks on three AA batteries. The stored readings can be read out over USB and plotted on a PC using gnuplot. Digital sensor modules keep the hardware simple and no calibration is required.

Kit of parts incl. PCB, controller, humidity sensor and air pressure sensor modules

Art.# 100888-73 • $50.20



Pico C Meter

Even upmarket digital multimeters boasting a built-in capacitance meter are useless if you want to check out tiny capacitances like 2.7 pF or 5.6 pF. Usually, you're tied to a lowest measuring range of 2000 pF, which is a good laugh to RF designers and radio amateurs. Although at 3.5 digits the DMM's resolution is 1 pF, any measurement below 200 pF or so will produce coarse if not ridiculous results. Pico C does a far better job. Beating many DMMs hands down, this little instrument easily measures capacitances down to fractions of a picofarad.

Kit of parts incl. Elektor Project Case, programmed microcontroller, LCD and PCB

Art.# 100823-71 • $118.40

Kits & Modules

ULTRA-LOW-LATENCY QUAD H.264 ENCODER

The **H264-ULL-SD4** is a high-performance, ultra-low-latency, quad-channel, H.264 encoder card on a single PCI/104 form factor board. The H264-ULL-SD4 provides a powerful and flexible solution for capturing and compressing up to four analog video inputs at full size and at full frame rate to the H.264/MPEG-4 AVC (Part 10) standard.

In addition to the ultra-low-latency full-frame rate encoding of four NTSC/PAL analog video sources, the H264-ULL-SD4 can also perform stream duplication to provide multiple encodings of the same input. This enables streams to be created at different resolution, compression settings dependent on requirements, and available bandwidth.

The H.264 video compression and bus utilization enables up to four H264-ULL-SD4 to be fitted in a PCI/104 system. This high-definition encoder card is well suited for applications such as defense, transportation, remote platforms, and situational awareness.

The H264-ULL-SD4 is supported by well-documented and established software development kits that minimize development risk and shorten time-to-market for applications requiring video recording or streaming.

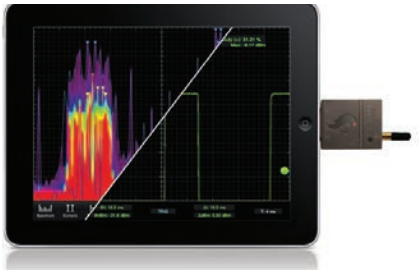
Contact Advanced Micro for pricing.



Advanced Micro Peripherals, Ltd.
www.ampltd.com

TOUCHSCREEN SPECTRUM ANALYZER, POWER METER, OR BOTH

Oscium recently announced a product line for the iOS test industry that enables an iPad, iPod touch, or iPhone to become a spectrum analyzer, a dynamic power meter, or both. Oscium's iMSO-104 merged a mixed-signal oscilloscope and the iOS family of products using the 30-pin dock connector. Oscium's new product line, **WiPry**, is the next installment in modular test equipment.



Three distinct products fit under the WiPry brand: WiPry-Spectrum, WiPry-Power, and WiPry-Combo (which combines the functionality of WiPry-Spectrum and WiPry-Power). WiPry-Spectrum leverages the colorful potential of the OpenGL interface on the iOS platform for real-time views of RF activity in the 2.4-GHz industrial, scientific, and medical (ISM) band. WiPry-Power crosses the chasm of this new platform by not only graphically displaying RF data from 100 to 2.7 GHz, but also adding the ability to capture, trigger, and record the actual power output of RF amplitude. An optional accessory kit is also available that enables users to make conducted measurements. WiPry-Combo combines all the features of both WiPry-Spectrum and WiPry-Power into one product.

WiPry is compatible with all generations of iPod touch, iPhone, and iPad devices running iOS version 3.1.3 or higher. It is made for iPod touch (1st, 2nd, 3rd, and 4th generation), iPhone 4, iPhone 3GS, iPhone 3G, iPhone, iPad 2, and iPad.

WiPry-Spectrum costs **$99.97**, WiPry-Power costs **$149.97**, and WiPry-Combo costs **$199.97**.

Oscium
www.oscium.com

DIGITAL DIFFERENTIAL PRESSURE SENSORS

The **SDP6x0-125Pa** and the **SDP6x0-25Pa** are two new versions of Sensirion's differential pressure sensors within the SDP600 series. The new sensors feature low measurement ranges of –125 to 125 Pa and –25 to 25 Pa, respectively. With their enhanced resolution, the sensors have improved zero-point accuracy of 0.1 Pa. In addition, zero-point repeatability reaches 0.05 Pa for the 125-Pa version and 0.03 Pa for the 25-Pa version. The SDP6x0-125Pa and the SDP6x0-25Pa are well suited for many applications in industrial and automotive markets as well as medical and HVAC markets where high-accuracy measurements of particularly low differential pressure are required.

In addition to the other products in the SDP600 series, the two sensors offer a digital I^2C output and are fully calibrated and temperature compensated. Thanks to the principle of calorimetric flow measurement, the CMOSens differential pressure sensors achieve exceptional sensitivity and accuracy, especially around the zero point.

Contact Sensirion for pricing.



Sensirion, Inc.
www.sensirion.com

Test Your EQ

Edited by David Tweed

ANSWERS for Issue 256

Answer 1—Yes, an "ideal" coaxial cable forms a perfect shield against the influence of external fields.

First, consider the fields generated by the signal in the first place. The electric field exists only between the center conductor and the surrounding shield, within the dielectric of the cable. Both the center conductor and the shield have magnetic fields associated with the currents flowing within them, but it's easy to show that the net magnetic field outside the shield is zero.

Thus, the signal in the coaxial cable produces neither type of field outside the shield. Applying the principle of reciprocity, it's impossible for an external field of either type to affect the signal.

Answer 2—"Ground" is in quotes for a reason—in any electrical network, only one point at a time can truly be considered the reference, or ground level. In other words, you can't really ground both ends of a piece of coax if they aren't located next to each other.

The point is, the "ground" connection provides a second path for the current that wants to flow through the shield of the coax, and this breaks the symmetry that enables the coax to protect against interference from external magnetic fields—or the outward leakage of the signal's magnetic field.

Answer 3—If the shield were a perfect conductor, currents produced in it by both internal and external fields would flow only along the corresponding surface and not penetrate its thickness at all. However, real cables are made with materials that have nonzero resistivity. This not only causes power loss (heating) within the cable, it also means that all currents penetrate the thickness of the shield to a greater or lesser extent, depending on frequency (skin effect).

Looking in detail at the electric field inside the coax, note that while it has radial symmetry, it is not purely radial at all points along the length of the cable. There is a longitudinal component that is due to the finite rate at which signals propagate through the cable. It is this component that is subject to leakage/interference because of an imperfect shield.

Other non-ideal characteristics of cables include gaps in braided shields, which can enable fields to "fringe" through the gaps. This effect gets worse as the frequency goes up and the gap becomes a larger fraction of the wavelength.

Answer 4—First of all, any sort of twisted pair implies the use of balanced signals, which means that the drivers and receivers are already designed to minimize the effects of common-mode signals. Also, all of the signal current flows in the twisted pair, and the shield does not carry signal current at all.

Secondly, twisted pair is inherently resistant to the effects of external fields (both electric and magnetic), because any effects (mostly) cancel out on adjacent half-twists. However, this effect backfires when the wavelength of the signal or interference approaches the pitch of the twist.

Contributed by David Tweed

What's your EQ?—The answers are posted at www.circuitcellar.com/eq/

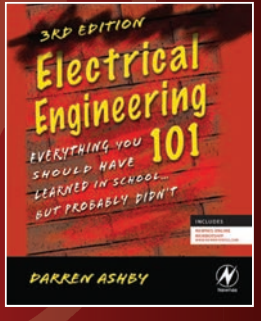
You may contact the quizmasters at eq@circuitcellar.com

Newnes Press

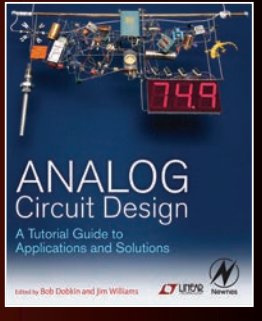
Education Never Ends. Everything you need to know to get started.



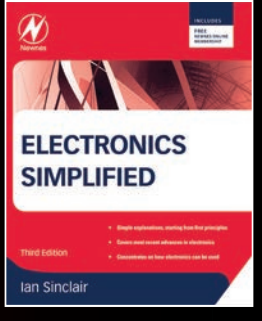
By: Keith Brindley
ISBN: 9780080969923



By: Darren Ashby
ISBN: 9780123860019



By: Bob Dobkin and Jim Williams
ISBN: 9780123851857



By: Ian Sinclair
ISBN: 9780080970639

Newnes

Save 20% on all titles when you purchase from www.newnespress.com.
Enter promotional code **Newnes30** at checkout.

Find us on Facebook!

Register for our e-news at newnespress.com
Receive our best discounts · Hear about books before they publish
Access to free sample chapters, video tutorials and more!

by Guido Ottaviani (Italy)

Sound Tone Detection with a PSoC (Part 2)

The Software & Sensing

The first part of this two-part series describes the hardware part of an audio sensor system. This article covers the design's software capabilities and how they can be used to measure the level of the signal, expand the dynamic of the amplifying chain, and reliably share this information in several different modes with the external world.

In the first part of this article series, I described the "hardware" capabilities of a Cypress Semiconductor PSoC 1 device and how to configure the modules inside to add a new sensor type to my Rino robotic platform (see Photo 1). The previous article shows how to multiplex three microphones on an amplifying chain and how to use a narrow band-pass filter to precisely detect a tone. A peak detector converts the signal to a DC level. Via an analog-to-digital converter (ADC), the software can read this level, act to optimize the programmable gain amplifiers (PGAs) to have the best dynamic, and send the measured level to the sensor board or whatever is waiting for this information with some different communication protocols.

Let's start with the software description.

SOFTWARE

After the analog blocks are connected, they are ready to condition the input signal to have a DC level proportional to the 4-kHz tone intensity. The digital blocks are ready to output the result in different ways. We now need some other glue to attain our goal. Some simple program code connects these two parts of the system returning the data we need.

The designer does most of the initializing job. When you have created the entire circuit, naming, connecting, and configuring all the blocks by selecting the "generate configuration files" menu item, the designer will create all the header, library, and include files needed, as well as the very basic main.c template to start writing your code.

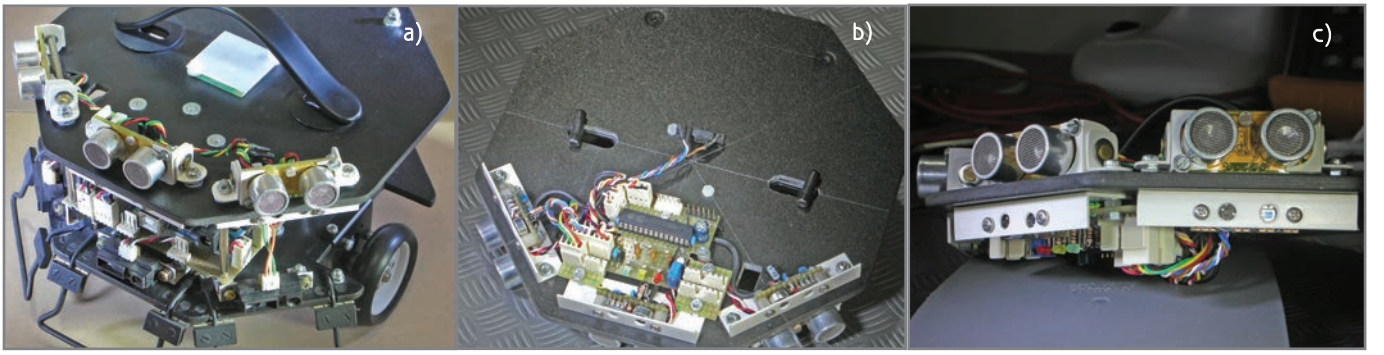


Photo 1a—An overall view of the Rino robotic platform with all its sensors: digital compass, ultrasound sonars, IR distance meters, and light and sound receivers. b—An upside-down view of the PSoC sound board installed among some other boards under the robot top cover. Maybe now it's clear why the platform is so strangely shaped. The board is connected to an I²C hub where all other sensors are connected. The white supports keep the microphones used by the sound board in place. c—Another view of the sensors installed on the robot top cover. Ultrasound sonars and one kind of light sensor are at the top, microphones and another kind of light sensor are under the cover.

If an interrupt is defined, it creates two library files for the module involved. For example, for the Timer module, it creates HB_Tmr.asm and HB_TmrINT.asm on which you must insert your code to start the interrupt service routine (ISR).

With the timer that generates an interrupt every 10 ms we have the "heartbeat" of the program. The ISR updates all the counters and flags used by other procedures to execute time-related actions.

READING THE ADC

This is another improvement. We don't have just a digital output that is only On/Off, but an analog output proportional to the input level that enables some more options with the software.

For each 10-ms timer cycle, the program performs one measure for the trimmer position. The three measures for the analog levels are performed alternatively, switching the three inputs every 30 ms. This is needed to enable the capacitor voltage to stabilize before readings. Every 10 cycles (i.e., every 100 ms) the program sends the values to the serial port.

AUTOMATIC GAIN CONTROL

By knowing the level of the signal at the output of the amplification chain and the gain of the two amplifiers, we know the actual level of the signal at the input. Therefore, we can adjust the gain of the amplifiers to avoid saturation of these stages (i.e., signal with a peak level too close to VDD). If the level goes above a given high threshold or below a low threshold, the gain will be adjusted to have maximum dynamics without saturation (see Figure 1).

SENDING TO THE UART

Using the library functions, it is easy to send data over the serial port. The high-level functions available are well described in the application programming interface (API) section of the datasheet. As explained in the datasheet, the API routines are part of the user module. Thus, you can deal with the module at a higher level. This section specifies the interface to each function together with related constants provided by the "include" files.

I've used the library functions a lot during the software development for debugging. It's easy to send the value of some variables of something similar to a communication program to monitor how the program acts.

After that, I'm using an I²C interface to communicate with the sensors board. But, even with a simple protocol, the same data can be exchanged through the UART. Refer to dsNav Communication, "Protocol Description," 2009 (www.guiott.com/Rino/CommandDescr/Protocol.htm).[1]

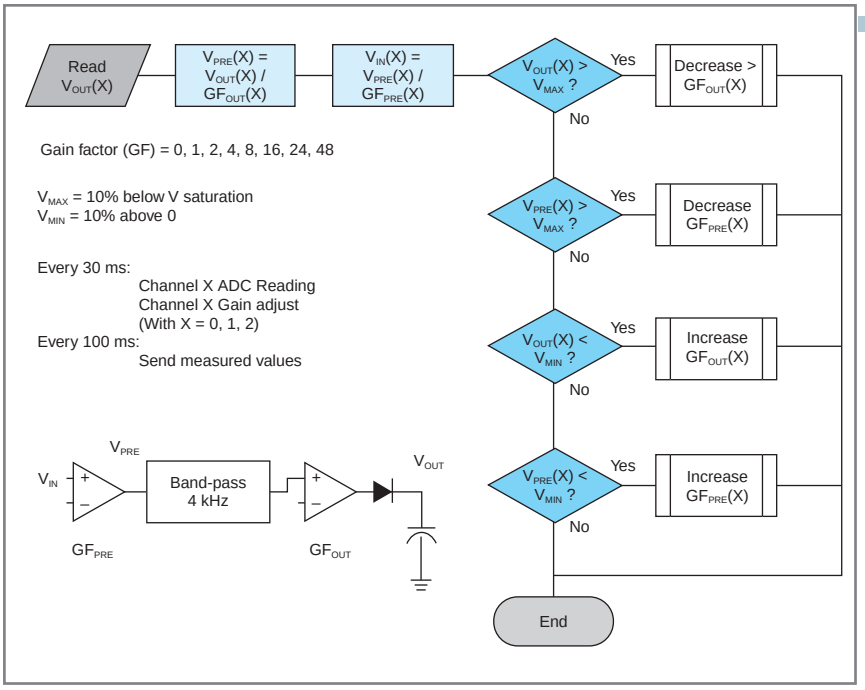


Figure 1—The procedure used for each of the multiplexed channels to adapt the gain of the two amplifier stages to the input sound level (AGC) in order to obtain the best dynamic and avoid saturation.

DRIVING LEDs & DIGITAL OUT

This part of the circuit is just used as backward compatibility with the old op-amp-based board. With the trimmer connected to an ADC port, you can manually set a threshold. When one input level goes over that threshold, the corresponding LED is switched on. The digital output port is set on if one or more LEDs are on.

USING I²C

At the beginning, I placed a standard I²C_HW module. This one is powerful and flexible with many options and slave, master, and multimaster capabilities. In a Cypress Semiconductor application note, there is a project for a bootloader on an I²C bus that becomes transparent to the circuit after program loading.

With all these features, it is not so simple to configure and manage. There are a lot of examples to study; but, in any case, you must still set and control many flags at some specific moments.

After a better analysis of my needs, I decided instead to use Cypress Semiconductor's EzI2C Slave 1.60 module, where "Ez" sounds like "easy." It really is easy, conciliating many people who are overwhelmed by the I²C management. It works only as a slave, with the same protocol of I²C EEPROM. Exactly what I need!

It is enough to place the module (no boxes occupied) and define a data structure that remaps the single-byte registers exposed to the I²C master:[2]

```
struct I2C_Struct { // I2C interface structure
  long I2C_MesValue;
} I2C_Regs
```

Initialize the module:

```
EzI2Cs_1_SetRamBuffer(sizeof(I2C_Regs), 0,
(BYTE *) &I2C_Regs)
```

And start it:

```
EzI2Cs_1_Start(); // Turn on EzI2C
```

In the initialization code, you must configure which registers are used, how many are read/write or read-only, and define a pointer to the data structure. That's all! The registers read and write transparently without any intervention.

The variable can be read by the master pointing at the specific register as a byte in the same way as an I²C EEPROM. Optionally, the ROM area can also be initialized with an equivalent syntax.

With the AGC, I expanded a lot the dynamics of the circuit, therefore "long" variables (4 bytes) must be used. This is exactly what is shown in the previous example. On the master side of the communication, you can use a similar structure to "repackage" the single-read bytes in 4-byte variables, but (there always is a but) you must be careful about the "Endianess" of your microcontroller. I'm using an Arduino board as a sensor board that collects all the interactions with the environment. It is the I²C master and it uses an Atmel microcontroller. This microcontroller uses a little Endian byte order: least-significant byte (LSB) first.

A PSoC uses a big Endian byte order: most-significant byte (MSB) first. Without an appropriate byte-order management you will have a reverted order with very strange values.

A valid alternative to the long variables is the usage of logarithmic values for the sound level instead of linear values. This is mostly used in the audio field and enables you to compress a long variable into a single byte, still maintaining a good resolution on low-level signal measurement. The dynamic range is 255 dB. It is difficult to design a circuit with so huge a dynamic. The range of this PSoC circuit, from the noise to the electret microphone saturation, is about 50 dB.

WATCHDOG

According to the I²C bus specifications, the slave must acknowledge the master exactly within the next clock cycle after the received message packet. To enable the use of I²C communication on a microcontroller that can't guarantee a deterministic response, the slave microcontroller can slow things down. The slave can hold the bus CLK line low until it is ready to send an ACK and reply to the master. This is the so-called "clock-stretching technique," and it hangs up the bus for all the time the microcontroller decides. Of course, the PSoC EzI2C module acts exactly in this way. It's very useful but very dangerous too. If the microcontroller, or its I²C module hangs for any reason, all the I²C devices on the bus will hang.

Because my robot must be reliable enough to perform a mission to Mars beside Spirit rover, I cannot accept this eventuality, even if it has a one in a million occurrence rate.

Enabling a watchdog on the PSoC, I'm sure this cannot happen. Or, if it occurs, the PSoC is reset, releasing the bus and starting from the beginning, missing just some measurement cycles.

The watchdog is enabled in the general preferences using designer IDE. It uses the same timer as the Sleep timer, dividing it by three (i.e., with an 8-Hz sleep timer, the watchdog resets the processor if not cleared for more than 375 ms).

To have complete supervision of the I²C bus, the master (sensors board) pats the dog to keep it awake by writing a specific slave register with a nonzero value every 100 ms. In the main 100-ms cycle of the PSoC program, this register is read. If the value is different from zero, the register and the watchdog timer (WDT) are cleared. If no WDT clear occurs for more than 375 ms—because of errors on the communication bus or because of an incorrect program behavior—the watchdog resets the microcontroller, resolving any problem.

INTERRUPT MANAGEMENT

Let me also criticize something a little bit. The PSoC is a microcontroller; therefore, it is possible to use interrupts in your program. This time too, the designer helps you to create the basic structure to manage the interrupt. But the entire process of assigning an interrupt to a module is not so easy and linear. Once defined, it works fine, but I've seen better ways to do it. In the hopes that this can help someone, I created a brief description of the operations to do in the exact sequence (see in Figure 2).

If you change the name of your block, you must restart everything from the beginning. A new file INT.asm is generated, and your code must be rewritten.

If you don't declare the ISR as an ISR with the `#pragma` directive, the compiler doesn't return any error, but the program will hang. The interrupt priority depends on

To assign an Interrupt Service Routine (ISR) to a block named e.g.: blk1, we must:

- Enable the global interrupt for the project using Designer configuration menu
- Place the blk1 module and enable its own interrupt management on properties windows, specific for each different kind of module
- Compile
- Look for ProjectName/lib/Library Source Files/folder in the workspace explorer window. After compilation, find the file "blk1INT.asm"
- Insert the following text after "Put your code here": "`ljmp _blk1_ISR_C`" exactly with this syntax, with every underscore included
- Come back to C code
- Enable global interrupts with the directive `M8C_EnableGInt`
- Insert the compiler directive: "`#pragma interrupt_handler blk1_ISR_C`" on main.c coded (without first underscore)
- Create your ISR. It must be named `blk1_ISR_C`
- End

Figure 2— A brief description of the operations (in the exact sequence)

ROUTE FASTER !



WITH PROTEUS PCB DESIGN

Our completely new manual router makes placing tracks quick and intuitive. During track placement the route will follow the mouse wherever possible and will intelligently move around obstacles while obeying the design rules.

All versions of Proteus also include an integrated world class shape based auto-router as standard.

PROTEUS DESIGN SUITE

Features:

- Hardware Accelerated Performance.
- Unique Thru-View™ Board Transparency.
- Over 35k Schematic & PCB library parts.
- Integrated Shape Based Auto-router.
- Flexible Design Rule Management.
- Polygonal and Split Power Plane Support.
- Board Autoplacement & Gateswap Optimiser.
- Direct CADCAM, ODB++, IDF & PDF Output.
- Integrated 3D Viewer with 3DS and DXF export.
- Mixed Mode SPICE Simulation Engine.
- Co-Simulation of PIC, AVR, 8051 and ARM7.
- Direct Technical Support at no additional cost.

Visit our website and use Promotional Code CLR2011JGB for an extra 10% Discount. Prices fromt just $249!

labcenter Electronics

www.labcenter.com

Labcenter Electronics Ltd. 411 Queen St. Suite 201, Newmarket, Ontario, Canada
Toll Free 866.499.8184, www.labcenter. com or Email: info@labcenter-electronics.com

THE NEW PICOSCOPE
2205 MSO
MIXED SIGNAL OSCILLOSCOPE
GREAT VALUE, PORTABLE, HIGH END FEATURES AS STANDARD AND EASY TO USE


Think Logically...
For your chance to WIN a PicoScope 2205 MSO visit www.picoscopemso.com and enter the code CC1
pico Technology

Channels	2 Analog, 16 Digital
Resolution	8 bit
Bandwidth Digital frequency	Analog 25 MHz, Digital 100MHz combined
Sampling rate	200MS/s
Trigger modes	Edge, Window, Pulse width, Window pulse width, Dropout, Window dropout, Interval, Runt pulse, Digital, Logic
Price	$575

www.picotech.com/pco464
1-800-591-2796

block position, according to a list that can be found in the 500-page technical reference manual. It could be easier!

BUILD A BOARD

You can now use this information to build an identical board if you need exactly the same behavior. Or you can use what you've learned to achieve your own goals. I hope this article series was useful to start using PSoC devices to develop similar circuits, or at least to "break the ice" with them.

Author's Note: The complete and up-to-date PSoC Designer project can be found as open source material on my Google Code space.[2]

Guido Ottaviani (guido@guiott.com) has made electronics and ham radio his hobby since the "tube times." He turned the hobby into a job working as an analog and digital developer for several years for an Italian communication company. Many years ago, a big change made him a technical manager at a company that develops and manages graphic, pre-press and press systems, and technologies for a large Italian editorial group that publishes sports newspapers and magazines. Some years ago, he got the scope and the soldering iron out of the drawer to make autonomous robots. Currently, Guido is an active member in various Italian robotics groups who shares his experiences with other self-professed "electronics addicts" and evangelizes amateur robotics.

PROJECT FILES

To download the "ADCINVRcalc.xls" spreadsheet for the calculation of parameters, go to ftp://ftp.circuit cellar.com/pub/Circuit_Cellar/2011/257.

REFERENCES

[1] dsNav Communication, "Protocol Description," 2009, www.guiott.com/Rino/CommandDescr/Protocol.htm.

[2] Google, Explorersound, http://code.google.com/p/explorersound.

RESOURCES

Cypress Semiconductor Corp., "PSoC 1 Overview," 2011, www.cypress.com/?id=1573.

———, Cypress PSoC 1 Application Notes Database, www.cypress.com/?app=search&searchType=keyword&keyword=&rtID=76&id=1573&applicationID=0&l=0.

G. Ottaviani, "Robot Navigation and Control (Part 1): Construct a Navigation Control Subsystem," *Circuit Cellar* 224, 2009.

———, "Robot Navigation and Control (Part 2): Software Development," *Circuit Cellar* 225, 2009.

———, "A Sensor System for Robotics Applications," *Circuit Cellar* 236, 2010.

Rino Robotic Platform, www.guiott.com/Rino/index.html.

SOURCES

Arduino board
Arduino | www.arduino.com

CY8C29x66, CY8C29466-24PX, CY8C27143-24PXI, AN1683, and AN2223 PSoCs, and EzI2C Slave 1.60 component
Cypress Semiconductor Corp. | www.cypress.com

HUMANDATA®

FPGA / CPLD Boards from JAPAN

SAVING TIME & COST with readily available FPGA boards

- Basic and simple features, single power supply operation
- Quality and reliability is provided by years of sales
- Same board size and connector layout – ACM/XCM series
- All stocked items are ready to be shipped immediately
- Over 100 varieties of FPGA/CPLD boards are available
- Customizing speed grade and/or any features are possible
- Free download technical documents before purchasing
- High quality and highly reliable FPGA/CPLD boards from Japan
- Almost all products are RoHS compliance

ALTERA FPGA Board

Cyclone IV E F484 FPGA board
ACM-023 series

Cyclone IV E | MRAM

EP4CE55F23C8N
EP4CE75F23C8N
EP4CE115F23C8N
Credit card size (86 x 54 mm)

RoHS compliant



CycloneIV GX F484 FPGA board
ACM-024 series

Cyclone IV GX | DDR2 | SIF40

EP4CGX50CF23C8N
EP4CGX75CF23C8N
EP4CGX110CF23C8N
EP4CGX150CF23C8N
Credit card size (86 x 54 mm)

RoHS compliant





Cyclone IV E F484 FPGA board
ACM-107 series

Cyclone IV E | MRAM

EP4CE55F23C8N
EP4CE75F23C8N
EP4CE115F23C8N
Compact size (43 x 54 mm)

RoHS compliant



Cyclone III USB-FPGA board
EDA-004

Cyclone III | USB Config. | USB Comm. | MRAM

EP3C55F780C8N
Credit card size (86 x 54 mm)

RoHS compliant



XILINX FPGA Board

Spartan-6 FGG484 FPGA board
XCM-018/018Z series

Spartan-6 | MRAM | DDR2

XC6SLX45-2FGG484C
XC6SLX75-2FGG484C
XC6SLX100-2FGG484C
XC6SLX150-2FGG484C
Credit card size (86 x 54 mm)

RoHS compliant



Spartan-6 FGG484 FPGA board
XCM-110/110Z series

Spartan-6 | MRAM | DDR2

XC6SLX45-2FGG484C
XC6SLX75-2FGG484C
XC6SLX100-2FGG484C
XC6SLX150-2FGG484C
Compact size (43 x 54 mm)

RoHS compliant



Virtex-5 FFG676 FPGA board
XCM-011 series

Virtex-5 | FRAM | SDRAM

XC5VLX30-1FFG676C
XC5VLX50-1FFG676C
XC5VLX85-1FFG676C
XC5VLX110-1FFG676C
Credit card size (86 x 54 mm)

RoHS compliant



Virtex-5 LXT FFG665 FPGA board
XCM-017 series

Virtex-5 | SDRAM | RocketIO | SIF40

XC5VLX30T-1FFG665C
XC5VLX50T-1FFG665C
Credit card size (86 x 54 mm)

RoHS compliant





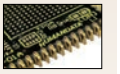
See all our products, download cables, A/D D/A conversion board, boards with USB chip from FTDI and accessories at :

www.hdl.co.jp/CC/

FPGA/CPLD Stamp Module PLCC68 Series

Easy and Quickly Mountable Module

FPGA Module PLCC socket mountable



- 50 I/Os (External clock inputs are available)
- 3.3V single power supply operation (Voltage converters for auxiliary power supply are built-in)
- Separated supply-inputs: Core, I/O drivers
- JTAG signal
- All PLCC68 series have common pin assignment
- Very small size (25.3 x 25.3 [mm])
- RoHS compliance
- MADE IN JAPAN



XILINX PLCC68 Series

XP68-01 Spartan-6 PLCC68 FPGA Module



Spartan-6 | PLCC 68

XC6SLX16-2CSG225C
16Mbit Configuration Device
Two User LEDs
One User Switch(Slide)

RoHS compliant

XP68-02 Spartan-3AN PLCC68 FPGA Module



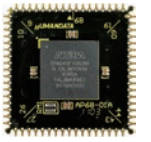
Spartan-3AN | PLCC 68

XC3S200AN-4FTG256C
FPGA internal configuration ROM
Two User LEDs

RoHS compliant

ALTERA PLCC68 Series

AP68-01 MAX II PLCC68 CPLD Module



MAX II | PLCC 68

EPM240F100C5N
EPM570F100C5N
External Clock inputs

RoHS compliant

AP68-02 MAX V PLCC68 CPLD Module



MAX V | PLCC 68

5M570ZF256C5N
External Clock inputs
On-board Voltage regulator

RoHS compliant

AP68-03 Cyclone III PLCC68 FPGA Module



Cyclone III | PLCC 68

EP3C10U256C8N
4Mbit Configuration Device
Two User LEDs
One User Switch(Slide)

RoHS compliant

HuMANDATA LTD. E-mail : s2@hdl.co.jp Fax :+81-72-620-2003

QUESTIONS & ANSWERS

IC Design Engineering and Beyond

An Interview with Monte Dalrymple



Monte Dalrymple has been designing integrated circuits (ICs) for more than 30 years. He holds a BSEE and an MSEE from the University of California at Berkeley and has 16 patents. Monte is the designer of all five generations of Rabbit microprocessors and author of the forthcoming book, *Microprocessor Design Using Verilog HDL.* In October, Monte and I talked about his various interests, from IC design to Verilog HDL.—*Nan Price, Associate Editor*

NAN: Where are you located?

MONTE: I live in Livermore, California, which is on the outskirts of the San Francisco Bay Area and just over the first row of hills from Silicon Valley. Livermore is home to the Lawrence Livermore National Laboratory, where they have the world's most powerful laser and several of the fastest computers in the world.

NAN: When did you first become interested in microprocessor design?

MONTE: I always knew I wanted to be an engineer, but it wasn't until I was asked to help out on a microprocessor project that I really became interested in designing central processing units (CPUs). I still consider myself an integrated circuit (IC) designer who has done microprocessors rather than a specialized microprocessor designer.

I seem to have passed the engineering gene on to my children. Three of them are practicing engineers (electrical, aerospace, and civil), and the fourth is studying mechanical engineering. Discussions at family get-togethers probably sound strange to non-engineers.

NAN: You have worked in the field of microprocessor design for many years. Tell us a bit about your background and experiences.

MONTE: In high school, I took vocational training in television repair. This was when all TVs used vacuum tubes and the voltages and hot components inside the cabinet could seriously ruin your day. That training was very practical, with almost no math involved, but it taught me invaluable debugging skills.

My undergraduate work at University of California at Berkeley was in "IC Design" that concentrated on analog bipolar technology. I think my only introduction to MOS technology at that time was in a device physics course. And, I only took one introductory logic design course as an undergraduate. But, the solid training in device physics and circuit theory has served me well over the years.

By contrast, my graduate work at Berkeley was almost entirely MOS technology, digital theory, and digital design. My Master's project investigated using switched-capacitor circuitry to extract the pitch information from speech. To guarantee reproducible results, I had to design and build a digital voice recorder using the just-introduced 4-Kb DRAM devices.

After graduation, I joined Zilog, Inc. as a design engineer, and my first project was the Serial Communications Controller (SCC) member of the Z8000 family. This design was done in 5-µm NMOS technology, making the transistors about 200 times larger than the current state of the art. All the schematics were hand drawn at the transistor level on sheets of paper vellum, and the layout was done using colored pencils on large sheets of gridded Mylar film.



Monte's workspace

There were no logic simulators in those days, so everything had to be verified by hand. Even the comparison between the schematic and the layout had to be done manually, with one engineer tracing the layout while another engineer followed along on the corresponding schematic. To make this verification robust, the original designer could not be involved, the engineer tracing the layout couldn't look at the schematic, and vice versa. I ended up helping to verify several of the Z8000 peripheral devices.

"One major project involved the manufacturer of an air data computer (ADC). An ADC is the part of an avionics suite that calculates airspeed and altitude based on sensor inputs. This is a critical component for flight safety, so the embedded software is subject to very stringent requirements by the FAA. The design used a military-specification Z8000 microprocessor that was no longer available, and the manufacturer of the ADC wanted to replace the Z8000 with an FPGA version so production could continue into the future."

After the SCC project was complete, I started work on a next-generation serial controller, but was quickly asked to help out on the Z800 microprocessor project. This design was supposed to be the successor to the Z80, but it was never really accepted by users. I designed the UART, refresh controller, cache controller, and memory controller. But when the designer of the execution unit left the company, I took over ownership of that block, also.

This was my first introduction to CPU design outside of a book, and it taught me a lot about how not to do it. The execution unit was a microcoded design with its own assembler, microcode memory, and a real microcode subroutine stack. But, an external memory read or write was done using a subroutine, with a one-clock cycle penalty on the entry to the subroutine and another one-clock penalty on the exit from the subroutine. So, the performance of the machine left something to be desired, and that's almost certainly why it never took off with users. This was when everything was switching over to CMOS, so being an NMOS design didn't help either.

I spent the next several years designing more serial controllers, culminating with the Universal Serial Controller (USC) family, where I was the architect and lead designer. This was my first CMOS design. By now, both the schematic entry and chip layout were being done on the computer, and logic simulation was available to verify the design.

My last major design at Zilog was the Z80380 microprocessor, where I was also the architect and lead designer. This design was what the Z800 should have been, with an efficient pipeline and no microcode. This was my first design using Verilog HDL, and it probably would not have been practical to do a design of this complexity using schematics.

NAN: In 1995, you started your own company, Systemyde. Tell us about your company. Can you describe some of the contract design work you do and provide a few examples?

MONTE: I started the company because I felt there was an opportunity to provide designs written in Verilog for use in custom ASICs or programmable logic. Regular semiconductor companies tightly control their intellectual property (IP), which makes it difficult for users when the standard chips don't have exactly the features they need or when a chip becomes obsolete. At the same time, I wanted to try to benefit a little more directly from any future designs I did, so Systemyde retains the core rights to nearly all the designs I have done since the company's inception.

I've described several of the projects I've done since starting Systemyde in *Circuit Cellar* articles, but there have been a number of other projects, also. For example, a major FPGA vendor contracted me to design a clone of the 8250 UART using the minimum number of configurable logic blocks. They wanted an outside designer to do the project to make sure that no special inside knowledge was employed in the design.

For another project, Zilog contracted Systemyde to get the Z80-SIO serial controller in Verilog format. They gave me access to the schematics so the design would be 100% compatible.

One major project involved the manufacturer of an air data computer (ADC). An ADC is the part of an avionics suite that calculates airspeed and altitude based on sensor inputs. This is a critical component for flight safety, so the embedded software is subject to very stringent requirements by the FAA. The design used a military-specification Z8000 microprocessor that was no longer available, and the manufacturer of the ADC wanted to replace the Z8000 with an FPGA version so production could continue into the future. The engineers on the project found the Systemyde website. I bid on the project and won the contract. It was much less expensive to put a clone of the Z8000 into an FPGA and qualify the new hardware than it would have been to rewrite and requalify the software. This was exactly the kind of project that I was anticipating when I started Systemyde, and now I am able to provide the Z8000 clone to other customers with similar requirements.

NAN: It's been 15 years since you wrote "Updating a Classic: the Z80380 Microprocessor" (*Circuit Cellar* 67, 1996). Tell us more about that project, and explain your involvement with the Rabbit series.

MONTE: I was the architect and lead designer for the

Z80380, which was a 16/32-bit backward-compatible upgrade to the original Z80 microprocessor. The design uses a 16-bit external bus, but is a 32-bit machine internally. As I mentioned before, this was what the Z800 probably should have been, and this was where I learned to design using Verilog.

The machine doesn't have the most efficient instruction set encoding because the marketing folks would not let me use any of the "No Operation" 1-byte opcodes to encode new instructions. In addition, the technology available for the project did not enable us to include any peripherals in the design. But the design is much faster than the original Z80, and Zilog stills sells the part.

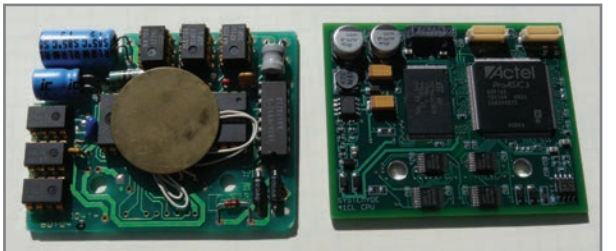
The Rabbit series of microprocessors is another example of where a customer needed features the manufacturer (Zilog) could not provide. Z-World, the company that became Rabbit Semiconductor, contacted me to help them design their own version of the Z180 microprocessor. I helped them define the architecture and did the entire design of the Rabbit 2000 myself. Seeing that they had a hit on their hands, we entered into a multi-year contract, and I designed the Rabbit 3000 and Rabbit 4000 for them.

Z-World was sold to Digi International, and I continued the relationship, leading the design of the Rabbit 5000 and the Rabbit 6000. These two generations used externally supplied IP, so I can't claim to have designed everything in them, but I did the CPU and all of the system integration for both of them.

NAN: In "Designing for Hostile Environments" (*Circuit Cellar* 210, 2008), you describe a CPU you built to orbit Jupiter. How did you become involved in this project and what exactly did you do?

MONTE: The University of Iowa has a long history of designing scientific instruments for spacecraft. In fact, James Van Allen, discoverer of the Van Allen radiation belts, was at the University of Iowa. The group there that designs these instruments has used the Z80 microprocessor as the main controller and has a large investment in software using that instruction set.

When they needed a Z80 design that could be inserted into a radiation-hard FPGA, they contacted me. Even though the FPGA they use employs triple-modular redundancy for each flip-flop, they wanted an additional level of protection. So, I modified my clone of the Z180, as I described in the article, and supported them in the effort to integrate the CPU with their custom logic. The Juno spacecraft that carries this scientific instrument was launched earlier this year and should arrive at Jupiter in July of 2016.



The CPU board used in Monte's calculator restoration. The original CPU board is on the left. The replacement CPU board is on the right.

NAN: You wrote a couple of articles geared toward taking the fear out of math: "Cyclic Redundancy Check Understanding CRC Without Math" (*Circuit Cellar* 193, 2006) and "Advanced Encryption Standard Understanding AES Without Math" (*Circuit Cellar* 235, 2010). Tell us about the "without math" theme.

MONTE: The mathematics underlying an algorithm like CRC or AES are obviously important for specifying the algorithm and for proving that it meets the design objectives. But I have found that the math can get in the way of figuring out how to implement an algorithm. As a designer, once I know that an algorithm works, I really only care about how to implement it and how to verify that I've done it correctly. With these articles I was trying to help other designers with those tasks.

NAN: In "Calculator Brain Transplant: Update a Vintage Design with FPGA Technology" (*Circuit Cellar* 243, 2010), you modernized a 30-year-old calculator with an Actel A3P125 ProASCI3 flash FPGA and a Xilinx XC2C64A 64 Macrocell CPLD. Tell us what inspired you to start this project. Have you made any updates to the calculator?

MONTE: I've always had a fondness for HP calculators. It was an HP-45 that got me through college, and I've used an HP-41CV throughout my career. When I found that HP had released the specifications for the custom microprocessor used in the 41C series, I had to try to implement the design in Verilog. Once I had that done, I realized that I could increase the performance and at the same time improve the capabilities of the CPU.

In the 41C series, the CPU, memory, and power supply are on a small daughterboard, so designing a plug-in replacement printed circuit board (PCB) was the obvious next step. Designing the board was a different challenge from what I am used to because I had to design the power supply, choose discrete components, and fit everything into a very small space.

Earlier this year, I sent beta-test boards out to users around the world in a dozen different countries. They found one hardware bug in the design, which was easily corrected in the FPGA, and suggested a couple of new features, which I also added. The main addition was the ability to substitute user code for the operating system in the calculator.

I haven't done any other updates yet, but I am thinking about doubling the size of the flash memory to 4 MB. At the same time, I am considering using a ferroelectric RAM instead of the static RAM to preserve the memory contents while the batteries are removed. There is one feature already in the FPGA that is not used in the design, which

is an interface for an SD card. Taking advantage of this feature will require a fair amount of software to be written to handle the SD card file format, but this will really bring the calculator into the 21st century.

NAN: In the near future, *Circuit Cellar* will be publishing your book, *Microprocessor Design Using Verilog HDL*. Simple question: Why Verilog HDL?

MONTE: Well, the obvious answer is because Verilog HDL is what I know and use. But Verilog HDL is much less verbose than VHDL. I think the cliché is "Write it once in Verilog or three times in VHDL."

NAN: What can readers expect to learn from the book?

MONTE: The book starts with an introduction to Verilog HDL, then follows along as I design a CPU. This is a real design, using the Z80 instruction set with an efficient two-clock machine cycle. I show how to implement the instruction set and cover the various tradeoffs I need to decide while implementing the design.

Most of the Verilog for the design is included in the book, and all of the information necessary for the reader to complete the design is present. I also include a testbench that can be used to verify the completed design.

My main goal in writing the book was to show that designing a microprocessor isn't really that difficult if you break the task into manageable pieces. While it's true that this is a fairly simple CPU by today's standards, the techniques I illustrate are applicable to designs of any size or complexity.

NAN: Any upcoming books or projects?

MONTE: I doubt there will be another book. Writing is painful for me because whenever I sit down to write I end up going back and editing what I have written previously. But I hope to continue doing articles for *Circuit Cellar*.

As far as new projects, I'm just finishing up a dual-channel HDLC controller that fits in an Actel A3P060 FPGA.

"My main goal in writing *Microprocessor Design Using Verilog HDL* was to show that designing a microprocessor isn't really that difficult if you break the task into manageable pieces. While it's true that this is a fairly simple CPU by today's standards, the techniques I illustrate are applicable to designs of any size or complexity."

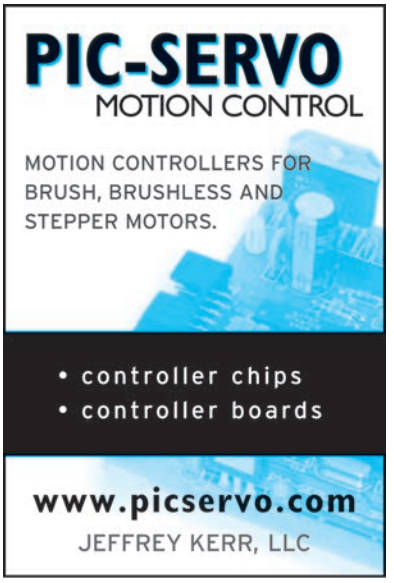
Using this small FPGA should actually be less expensive than using one of the dedicated serial controllers I've designed in the past. Beyond that, I'm always looking for new projects.

NAN: What types of technologies do you think will be important in the near future?

MONTE: I was initially very reluctant to switch from schematics to Verilog HDL. But now I would never go back to schematics, because Verilog enables me to be so much more productive. I expect that any technology that enables even more design abstraction, and hence more productivity, will be important in the future.

If my calculator project taught me anything, it is the importance of low-power operation. I tried to squeeze every microamp out of the standby power in that project. Anything that reduces the overall power consumption of a design is going to be important in the future.

As far as hardware, I think hardware design using programmable logic is going to become much more common. Just as with microprocessors, the newest, fastest, or biggest FPGA gets all the press coverage, but I think there are lots of places in embedded systems where small FPGAs make a lot of sense.





by Curt Terwilliger (USA)

The (VI)sualizer

A Smart Electronic Load

The (VI)sualizer is a lab instrument that enables you to examine an energy source's voltage (V) and current (I) delivery ability. An MCU accepts input data from local and remote interfaces, commands the desired load current, and monitors operating conditions.

The "(VI)sualizer" is a lab instrument for profiling solar, chemical, and grid-powered energy delivery devices (see Photo 1). It gets its name from the fact that it enables the user to examine—or visualize—the voltage (V) and current (I) delivery ability of an energy source when strenuous loads are applied. It is, in essence, a smart electronic load.

The (VI)sualizer enables a user to draw a test current from an energy source and then see to what degree changes in current affect the test voltage. From this information, it can calculate parameters such as the load resistance into which a solar cell delivers its maximum amount of energy, the amp-hour (Ah) capacity of a battery, or, using pulsatile loads, the voltage-regulation accuracy of a conventional power supply.

My goal was to produce a design that would include signal generation, testing algorithms, and data capture/storage in a single instrument—needing nothing more than a connection to a web browser to make a complete data analysis system. This would be an advantage over lash-ups using several pieces of commercially available gear to achieve the same results.



Photo 1—The main operating controls are (counter-clockwise from the input jacks): the Load Type selector, which sets the operating mode; the coarse and fine load-current controls; the panel meter and its associated Function switch; and the pulse-control group, which sets pulse width, duty cycle, and amplitude.

DESIGN SUMMARY

The block diagram shown in Figure 1 illustrates the main components of the (VI)sualizer. The heart of the system is an NXP Semiconductors LPC1768 mbed module, which accepts input data from the local and remote interfaces, commands the desired load current, and executes the tests.

The load current is controlled by a bank of power MOSFETs mounted on fan-cooled heatsinks, controlled by a digital-to-analog converter (DAC)/op-amp drive circuit that is used to set the target current. The DACs are driven through an optical isolation interface. A voltage-controlled oscillator is used to sense the applied voltage. The oscillator's output is also optically isolated.

Two storage devices have been added to the mbed's native storage. The first is a 4-GB flash memory in the form of a USB flash drive, which is used to store test data files. The second is a small EEPROM attached by I^2C and used for nonvolatile storage of set-up data. This data is automatically saved when a power failure is detected by means of an interrupt input driven from the AC line.

A metering circuit shows operating conditions, including three current ranges, a voltage scale, a wattage scale, and a pulse-width/duty-cycle scale. The wattage and pulse/duty signals are generated by pulse-width modulation (PWM) outputs on the mbed; the others are generated directly in the hardware. Another PWM signal drives the variable-speed cooling fans according to the temperature reading reported by a sensor attached to the heatsink.

SAFE & CURRENT MODES

The user selects the operating mode of the instrument with the Load Type control, which is located above the input jacks on the right side of the front panel (see Photo 1). If Safe mode is selected, the input current is set to zero, so the load is effectively disconnected. If the Voltage or Power modes are selected, the instrument control is enabled via a connected web browser. More on this later. First, let's look at the Current mode, which is used to test power supplies.

When Current mode is selected, the Load Value control sets the current drawn from the input, using the outer knob to control the coarse value and the inner knob to control the fine value. The mbed's analog-to-digital converter (ADC) reads these potentiometers and software shaping is applied to the values. Thus, the range from 0 to 5 takes about the same amount of rotation as the range from 10 to 20. The fine control acts as a vernier dial, varying the setting of the coarse control by ±5%, again by means of software shaping.

If the coarse setting is below a minimum value (about 1% of full scale), then the fine setting directly sets the current from 0 to 200 mA. This switchover makes it possible to set extremely small currents with a precision that could not be obtained using the coarse control/fine vernier arrangement described earlier.

It is also possible to introduce pulsed currents using the four pulse-related controls and connectors grouped in the lower left section of the panel. Pulse width and duty cycle, for example, are set by a rotary encoder. The encoder's integrated push switch is used to change the role of the encoder and to announce its current role with the pulse or duty cycle LEDs. These LEDs blink, illuminate steadily, or are extinguished according to the operating mode selected.

The pulse width may be set from 1 ms to 10 s. This 10,000:1 span is made user-friendly by a speed-sensitive encoder algorithm. A normal rotation of the control will change the setting by 20 counts per revolution, a swifter rotation will advance the value by a factor of about two, and a real twist of the knob will move the output value by an order of magnitude or so.

The value of current drawn by the (VI)sualizer during a pulse is set using the coarse/fine concentric controls in exactly the same manner as the Load Value controls. A three-position toggle switch turns the load pulses on and off. In addition, the switch may be placed in the Set position to enable exact setting of the pulse amplitude. The selected width or duty cycle LED will blink to remind the user that the instrument is in Set-up mode. A transistor-transistor logic (TTL)-level synchronizing signal is provided at the front panel sync out jack so an oscilloscope can be triggered when the (VI)sualizer is about to generate a pulse.

BROWSER-CONTROLLED MODES

The Voltage mode is selected when the user wants to measure the current-time product available from an energy source whose voltage decays as energy is withdrawn—typically a rechargeable battery. Its user interface is through the Battery Test tab of the web interface. Refer to the *Circuit Cellar* FTP site for an illustration of this page.

Key inputs are the expected milliamp-hour (mAh) capacity

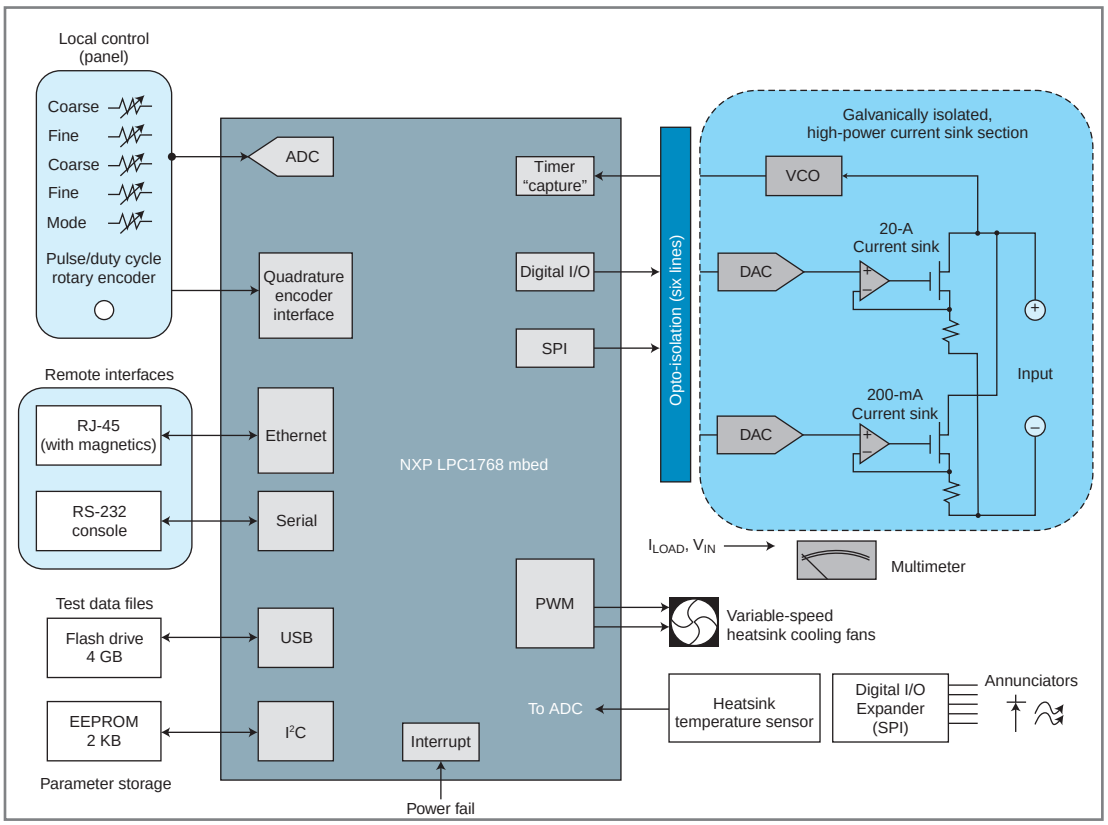


Figure 1—The mbed module accepts input commands from the local and/or remote interfaces, controls the desired load current in accordance with the selected testing algorithm, and stores or displays the resulting data.

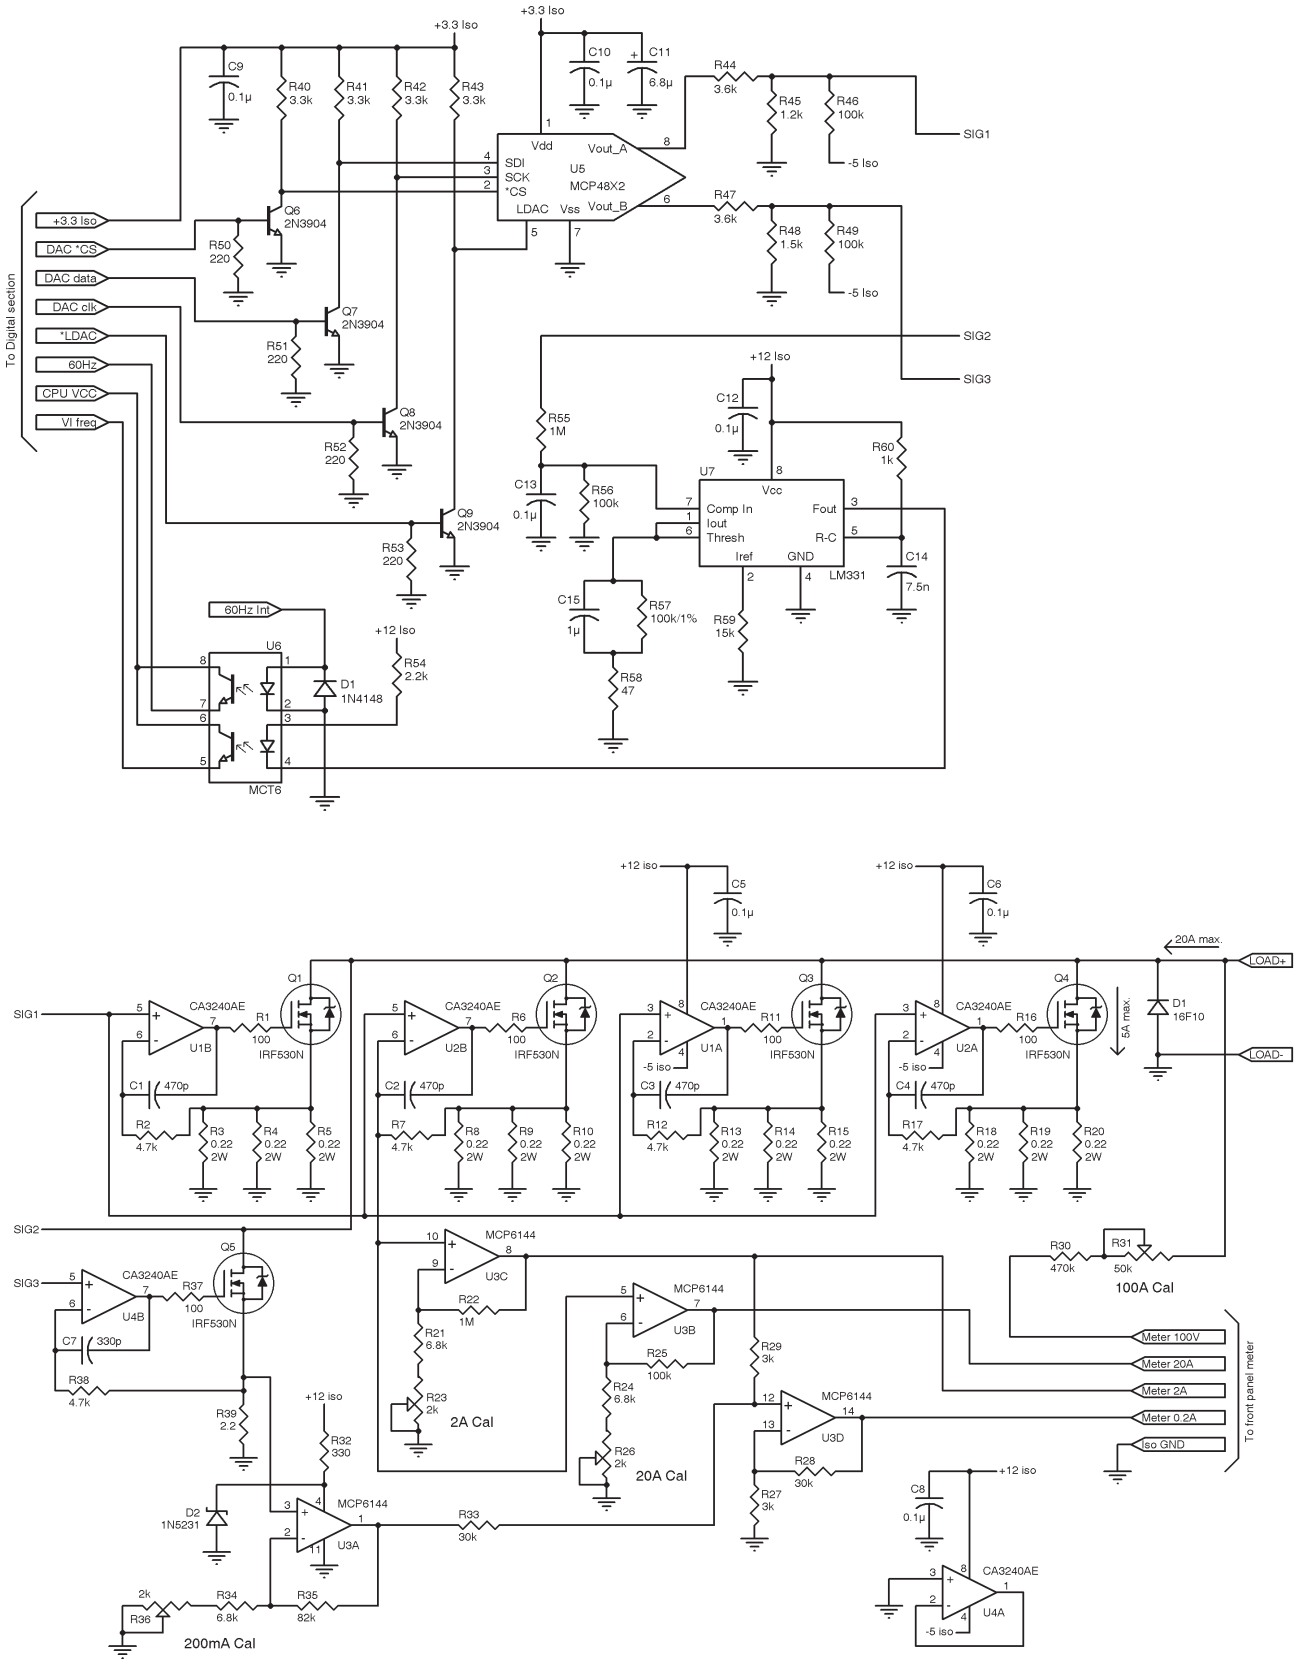


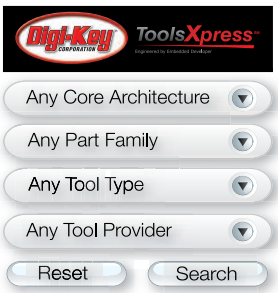
Figure 2—The analog portion of the circuit comprises an optically isolated dual DAC, five high-power current sink sections, meter-driving circuits, and an optically isolated VCO for voltage sensing.

www.DevtoolsXpress.com

Pick a Tool. Any Tool.

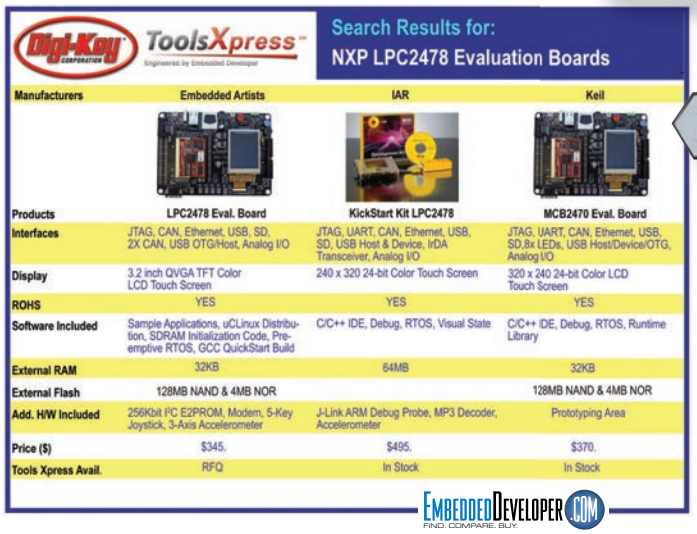
Find the Right Development Tool, Compare it to Other Tools, Evaluate It, and Buy It from Digi-Key Tools Xpress -- Without Leaving Our Site.

The Digi-Key Tools Xpress intuitive research engines are used by engineers worldwide to locate, compare and evaluate hardware or software development tools.



Digi-Key Tools Xpress, engineered by Embedded Developer, is the only site in the industry where engineers can quickly find, compare and buy the leading development tools.

Compare before you buy: tools are listed side-by-side, with features and performance specs, availability, and prices, so you can make an educated decision!


Digi-Key ToolsXpress — Search Results for: NXP LPC2478 Evaluation Boards

Manufacturers	Embedded Artists	IAR	Keil
Products	LPC2478 Eval. Board	KickStart Kit LPC2478	MCB2470 Eval. Board
Interfaces	JTAG, CAN, Ethernet, USB, SD, 2X CAN, USB OTG/Host, Analog I/O	JTAG, UART, CAN, Ethernet, USB, SD, USB Host & Device, IrDA Transceiver, Analog I/O	JTAG, UART, CAN, Ethernet, USB, SD,8x LEDs, USB Host/Device/OTG, Analog I/O
Display	3.2 inch QVGA TFT Color LCD Touch Screen	240 x 320 24-bit Color Touch Screen	320 x 240 24-bit Color LCD Touch Screen
ROHS	YES	YES	YES
Software Included	Sample Applications, uCLinux Distribution, SDRAM Initialization Code, Preemptive RTOS, GCC QuickStart Build	C/C++ IDE, Debug, RTOS, Visual State	C/C++ IDE, Debug, RTOS, Runtime Library
External RAM	32KB	64MB	32KB
External Flash	128MB NAND & 4MB NOR		128MB NAND & 4MB NOR
Add. H/W Included	256Kbit I²C E2PROM, Modem, 5-Key Joystick, 3-Axis Accelerometer	J-Link ARM Debug Probe, MP3 Decoder, Accelerometer	Prototyping Area
Price ($)	$345.	$495.	$370.
Tools Xpress Avail.	RFQ	In Stock	In Stock

EmbeddedDeveloper.com FIND. COMPARE. BUY.



Join the thousands of engineers worldwide who use Digi-Key Tools Xpress for their development tool needs.

Digi-Key CORPORATION

ToolsXpress℠
Engineered by Embedded Developer

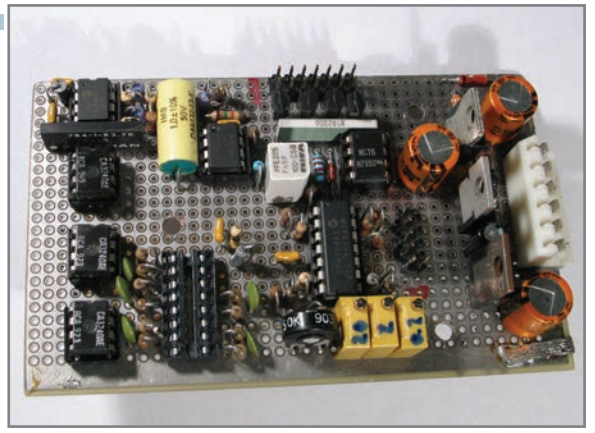


Photo 2—The analog board was constructed on perf board and mounted on insulated standoffs to maintain galvanic isolation. The 14-pin stake header at the top connects the optical isolation circuits on this board with those of the digital board. The right edge of the board contains the 12-, 3.3-, and –5-V regulators for the isolated power supplies.

of the battery to be tested, the test current to be used, and the voltage at which the test is to be cut off. All three values are given bounds checks and are automatically restored to "sane" values if an erroneous entry is made. After this data is entered, the user clicks Start Test and the test runs to completion. Results are posted to the webpage at the end of the test.

Similarly, the Power mode is selected in order to find the optimal load and power output for an energy source with non-negligible (and possibly nonlinear) source resistance—typically a solar cell. This test is totally automated. The user just clicks Start and the (VI)sualizer automatically ramps up the test current, calculates the power being delivered at each new test-current value, and reports results once the maximum power is reached.

HARDWARE IMPLEMENTATION

I began the design with the DAC and analog sections. I chose a Microchip Technology MCP4812 DAC, which is a SPI-controlled, dual-channel, 10-bit device with a 0-to-2.047 output voltage range. One channel was used to drive the 20-A current sink, the other channel was used to drive the 200-mA current sink.

Five current-sink modules are used in the load. Each comprises a power MOSFET and a fan-cooled heatsink that can support 25 W of dissipation. I used an op-amp section to individually drive each MOSFET. This ensures optimum current sharing. While MOSFETS cooperate beautifully in saturated switching applications, showing a negative coefficient of drain current versus temperature, the reverse is true when they are operated in the linear region.[1]

The input capacitance of each MOSFET is on the order of several hundred picofarads (pF), which is a tricky load for the op-amps to drive. I found it necessary to add an outboard compensation network in order to tame the tendency to oscillate.

In order to enable the load to be used in both positive-ground and negative-ground applications, the "ground" reference of the current sinks is galvanically isolated with optocouplers.

To keep the optocouplers switching as quickly as possible, I avoided running them into heavy saturation. Instead, their quasi-linear outputs drive common emitter transistor switches that do not slow down too much when they saturate. This enabled me to achieve usable transmission frequencies of 70 kHz or so.

I used a National Semiconductor LM331 precision voltage-to-frequency converter to sense applied voltage and generate a variable-frequency signal suitable for transmission across an optocoupler interface. Given the optocoupler's speed limitations, I chose a maximum VCO frequency of around 70 kHz, which would be produced at the maximum input voltage (corresponding to 100 V).

Photo 2 shows the physical implementation of the analog circuits. Figure 2 is the schematic.

THE CPU

Much of the digital circuitry needed is already present on the mbed board. I added a few parts, as shown in the digital board schematic, to support additional memory and I/O (see Figure 3). The digital section includes the mbed module, a USB flash-memory stick, and an RJ-45 jack used to connect to the Ethernet. The remaining chips are the shift register for the SPI port expander, a Maxim Integrated Products MAX232 line driver/receiver, the EEPROM parameter memory, and two optical isolators.

The mbed module itself was modified to bring out the PhA and PhB signals from the internal quadrature encoder interface (QEI) module.[2] Small-gauge wire was used for the connections (see Photo 3).

SOFTWARE IMPLEMENTATION

The program's main modules are an

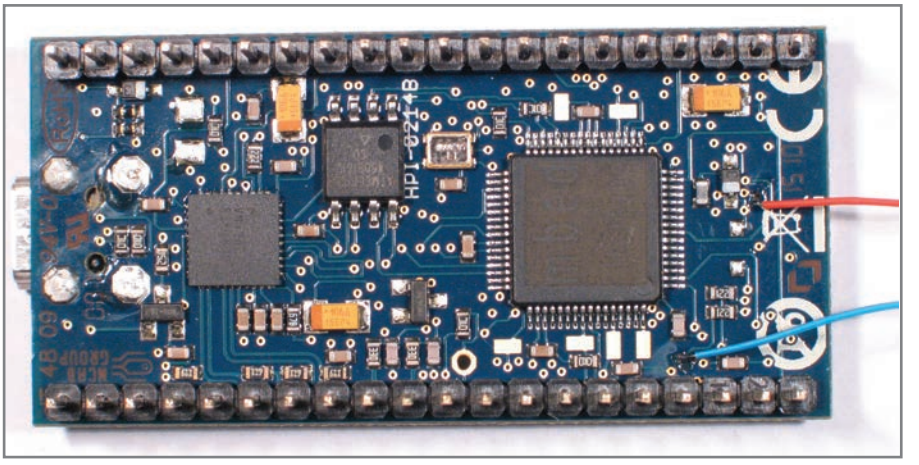
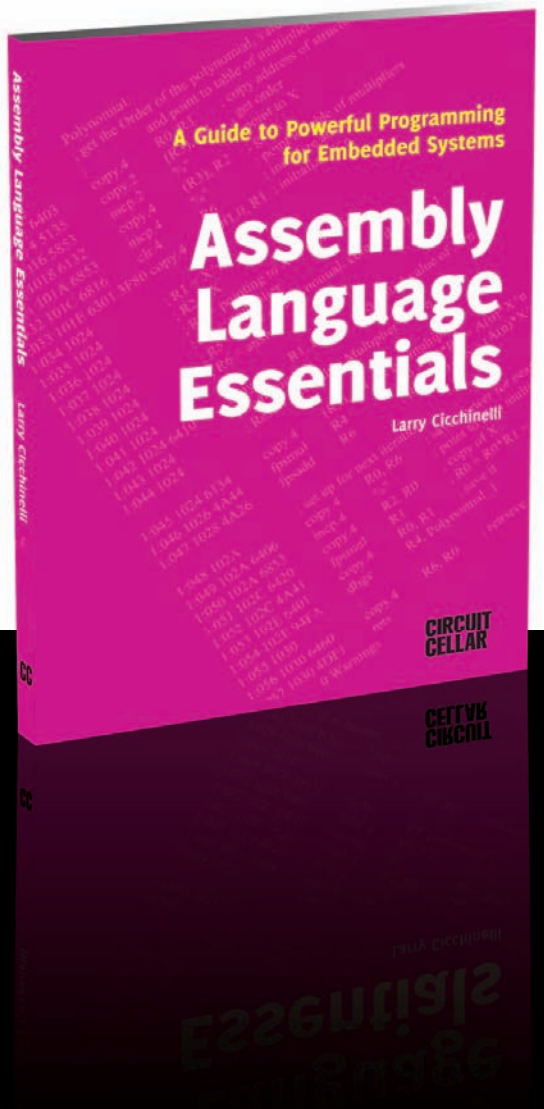


Photo 3—The mbed module was modified slightly to bring out the PhA and PhB signals from the internal QEI module. Rather than soldering directly to the pins of the NXP LPC1768 chip, wires were connected to feed through vias in the traces running from the pins of interest to their respective LED dropping resistors (R20 and R22 on the mbed).

Assembly Language Essentials

Circuit Cellar's first book, ***Assembly Language Essentials***, is a matter-of-fact guide to Assembly that will introduce you to the most fundamental programming language of a processor.



Author Larry Cicchinelli provides readers with:

- An introduction to Assembly language & its functionality
- Essential terminology pertaining to higher-level programming languages & computer architecture
- Important algorithms that may be built into high-level languages — multiplication, division, and polynomial evaluation
- Overview of Interrupt Service Routines
- Free, downloadable Assembler program …

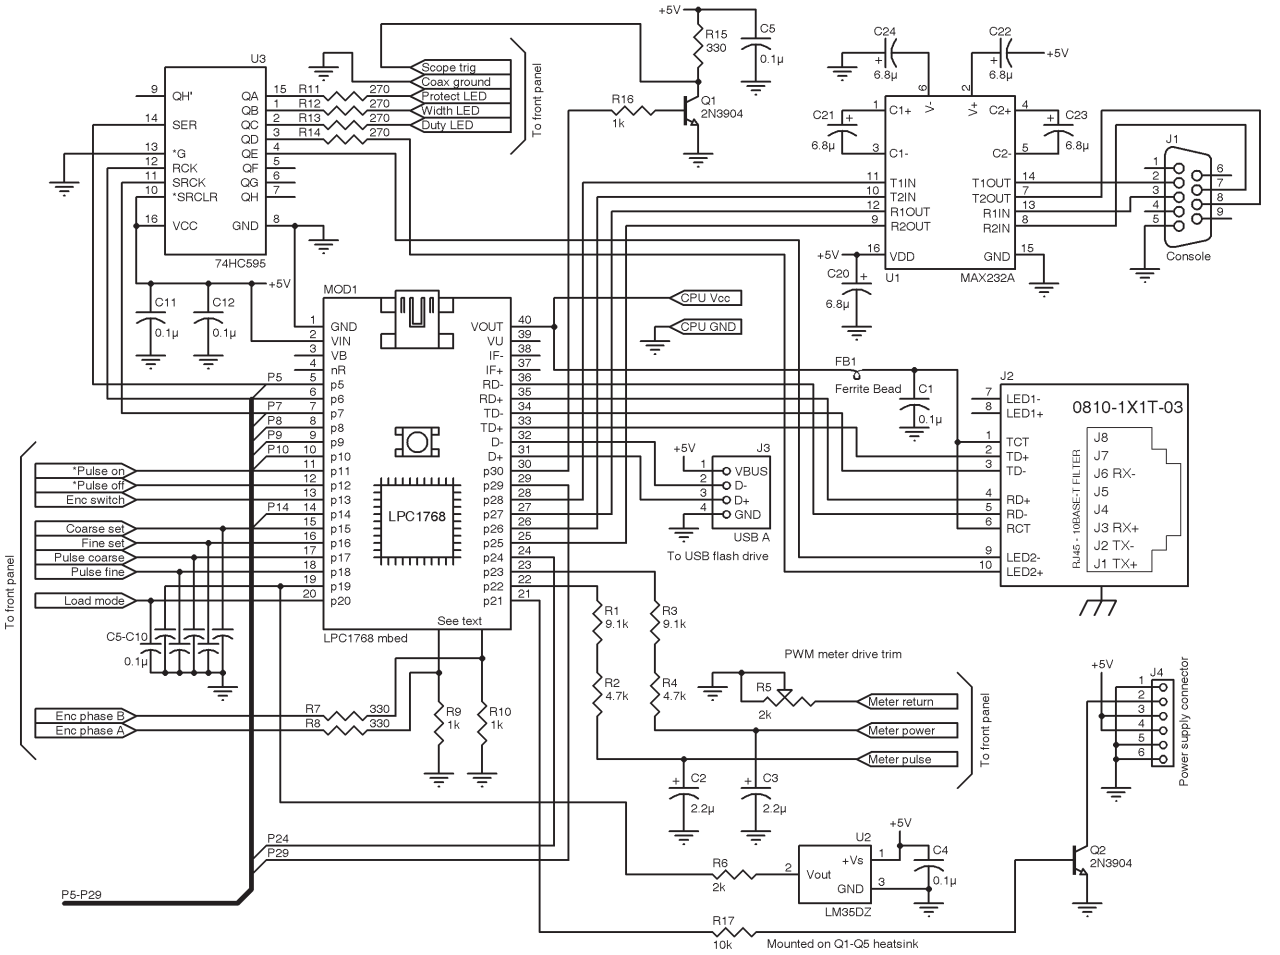
and more!

On Sale NOW!

$47.50

WWW.CIRCUITCELLAR.COM • WWW.CC-WEBSHOP.COM

CIRCUIT CELLAR
THE WORLD'S SOURCE FOR EMBEDDED ELECTRONICS ENGINEERING INFORMATION



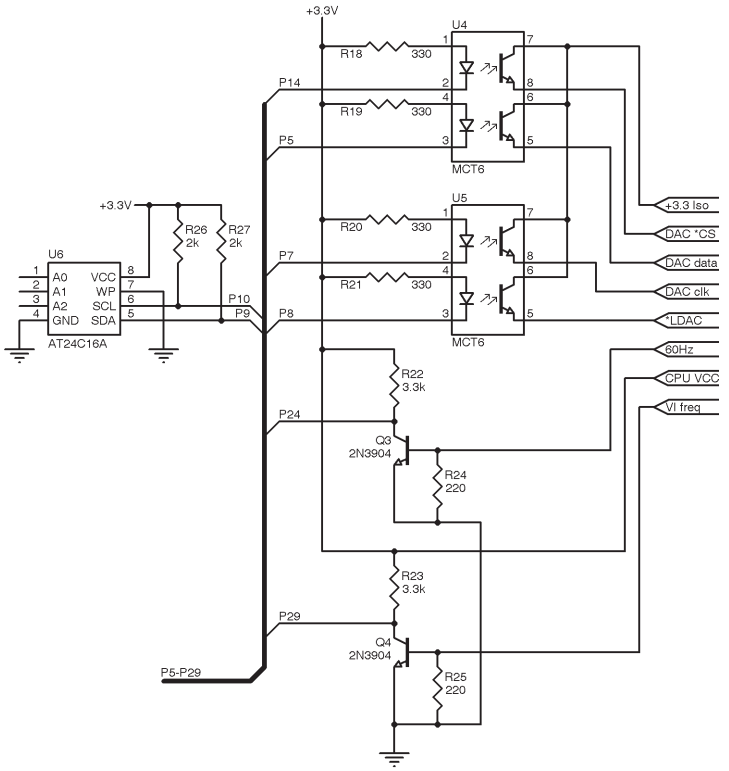


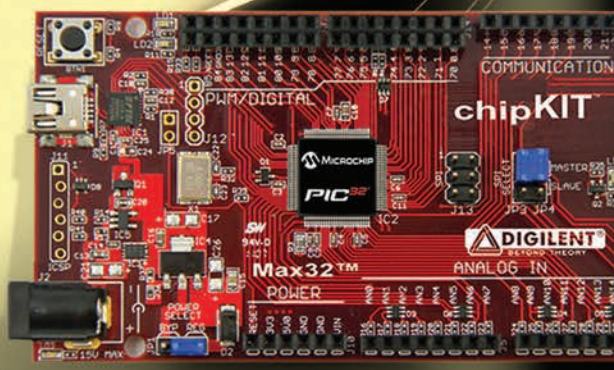
Figure 3—The digital portion of the circuit is composed of the mbed module with its 32-bit Cortex M3 CPU, the NXP LPC1768. The (VI)sualizer makes use of the mbed's Ethernet, serial, I²C, SPI, ADC, and PWM modules. For details on the power supply and front panel interconnections, refer to the *Circuit Cellar* FTP site.

DESIGNSPARK

Turn a hot idea into a cool solution.

DesignSpark chipKIT™ Challenge

DesignSpark is the perfect environment for engineers to work together, challenge each other to think outside the box and engage with fellow designers from around the world. Participants of the DesignSpark chipKIT™ Challenge are encouraged to post photos/videos, summaries, and questions about their projects, code, and PCB designs at **www.chipkitchallenge.com**. Register today and you could be selected for a chipKIT™ Community Choice Award!



Prizes include a $100 voucher for RS Components/Allied Electronics and a free digital subscription to *Circuit Cellar* and *Elektor* magazines. Community Choice Awards will include a winner of the most creative project proposal, selected by the DesignSpark team, and most popular project, selected by the DesignSpark community.*

Visit www.chipkitchallenge.com

for complete rules and details.

* Participation in the Community Choice Awards does not increase your chances of winning the Grand Prize with your Final Project(s) submission. See website for more information.

IN ASSOCIATION WITH:

elektor · MICROCHIP · RS · CIRCUIT CELLAR · ALLIED ELECTRONICS · DIGILENT BEYOND THEORY

DesignSpark chipKIT™ Challenge

chipKIT™ is a registered trademark of Microchip Technology Inc. Max32™ is a registered trademark of Digilent, Inc.

initialization section, a main task loop, a collection of task modules, and the interrupt handlers. Also, associated webpages were developed for this project and are served by the mbed. Finally, there are two new drivers written especially for this project.

The initialization module loads configuration parameters from the EEPROM or from console inputs if the user has requested Set-up mode during boot-up. If an Ethernet link is detected, the module initializes the stack, fetches time data from a network time protocol (NTP) server, and, if so configured, sends a tweet with the current IP address. Following these steps, the main task loop is entered.

The main task loop executes tasks sequentially. Each task is designed to be nonblocking. The main elements of the task loop are shown in Figure 4. Descriptions follow.

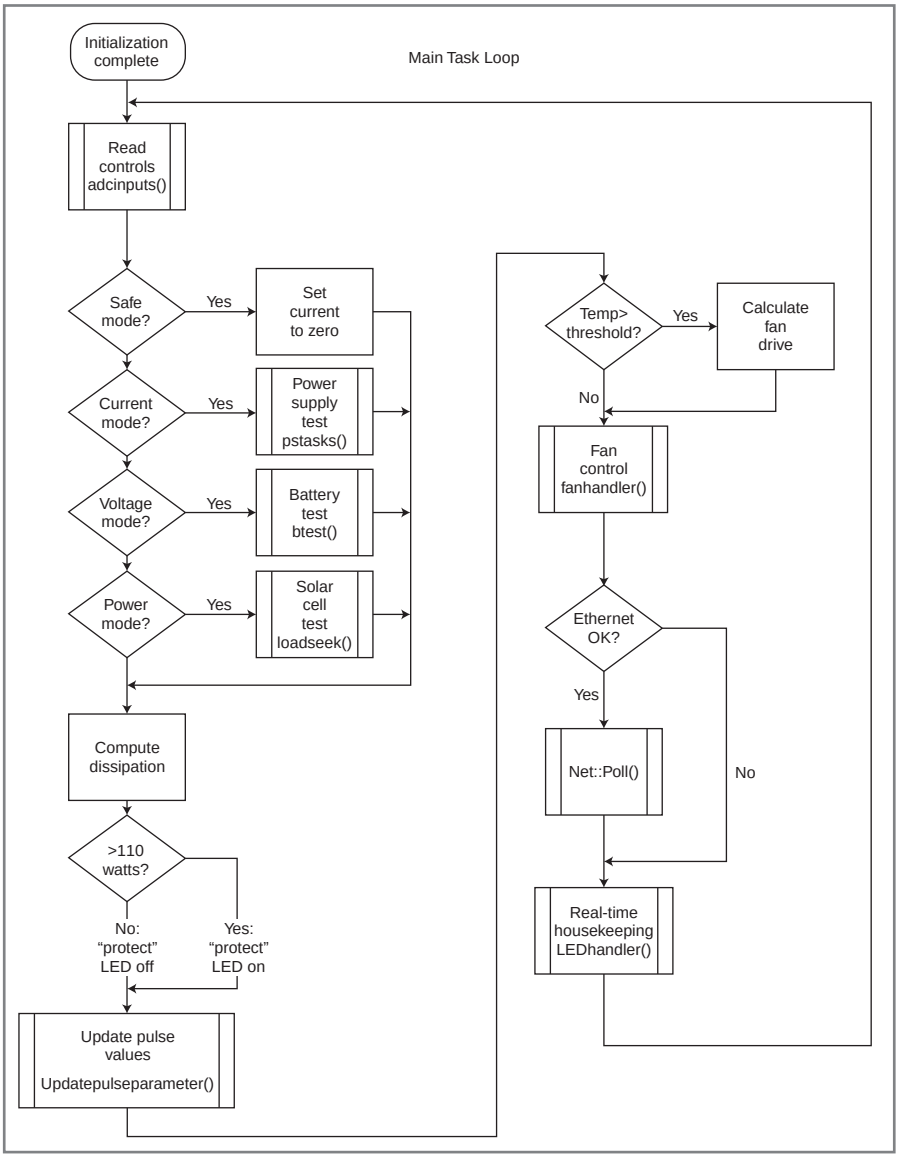


Figure 4—A flowchart of the main task loop. Testing algorithms are selected according to the operating mode. Housekeeping tasks look after real-time chores including Ethernet polling, front-panel control settings, LED handling, and fan operation. The numerous state machines are designed to be nonblocking, so that none get starved for CPU time.

SAMPLING THE ANALOG INPUTS

The first module called is `adcinputs()`, a state machine that reads the controls and sensors using the ADC's six channels. The handler collects results from the ADC, deglitches them, and averages the results. Each time the handler is called, it checks to see if the ADC has completed its most recent conversion. If not, the handler simply returns control to the main loop so no blocking occurs.

If a conversion has completed, it is compared to the last data for the channel under consideration. If the two data values match (to within a specified tolerance), then no "glitch" has occurred. This means the data is valid, so it is used to update the average value for the channel.

DETERMINING THE MODE

Next comes a set of tests to determine the operating mode according to the load-type switching (which was sampled by the ADC module). If in Safe mode, the current sinks are set to zero and control moves to the next portion of the task loop.

THE CURRENT MODE

If in Current mode, the `pstasks()` module is called. This module has the job of handling the raw input data from the front-panel load and pulse controls. That data is reported as an analog level between 0 and 4,095, which is put through a piece-wise linear remapping to give the controls the "spread" dial previously mentioned.

Following that remapping, the data is converted to DAC counts that deliver the desired output current. For example, the DAC driving the 20-A current sink section is sent a value equal to:

$$\text{DACdata} = \left(\frac{\text{mapped control value}}{1.172}\right) + 390$$

This calculation is carried out for both high- and low-current DAC sections and for both load current and pulse current.

This module also calculates an "effective current," which combines the load current and pulse current according to the current pulse-duty cycle. This value is used to determine actual power dissipation for the protection and fan modules.

THE VOLTAGE MODE

If the load-type switch is set to Voltage, the `btest()` state machine is called. This finds the amp-hour capacity of the connected battery.

Once triggered by the user (through an RPC variable handled by JavaScript on the webpage), the state machine begins a sequence of measuring and logging data.

The process begins by starting a ticker that will be used to mark sample times. The ticker interval is calculated in order to produce 1,000 samples (assuming the battery delivers the milliamp-hour-energy specified by the user). Following this, a new file is created on the USB file system and the test parameters are written to it—expected milliamp-hour, test current, cutoff voltage, and current time and date. The file is then closed and the state machine moves to the WAIT state. This routine is nonblocking, so it returns control to the main task loop after every action. When each sample interval expires, the state machine moves to the SAMPLE state, where current time and voltage are appended to the data file. Then, the voltage is checked to see if it has reached the cutoff point. If so, the test is ended. If not, the test continues.

Note that the data file is opened, written, and closed for each sample. I chose to do this rather than leave the file open over extended periods of time. I thought this would make things more robust in case of an unexpected power failure or similar event.

THE POWER MODE

The final possible position for the load-type switch is the Power mode. Another state machine is used to implement the `loadseek()` module, which finds the optimum load for a solar cell or other high-impedance source.

The test is initiated by the user clicking a button on the related webpage. That action is then relayed by a JavaScript-managed RPC variable. The state machine sets the output current to zero, then starts a timer and moves to the WAIT state. The timer interval is defined in the program as 250 ms.

When the timer expires, the state machine next moves to TEST_POWER, which calculates the power generated at the present setting. (This calculation happens at the end of the WAIT period so the source has time to stabilize after each current step.) If the power is less than the previous power, the test has found the maximum power point. In that case, the results are posted and the test loops back to the start.

If a match has not been found, the state machine moves to the RAISE_CURRENT state. This increases the output current by an amount that depends on its present value. If the present value is less than 220 mA, steps of 2.2 mA are used. If the value is between 0.22 and 2.2 A, steps of 27 mA are used; otherwise, steps of 54 mA are used. (The non-round numbers are artifacts of the hardware design. They reflect the natural step size for the DAC/MOSFET combinations.)

The cycle of raising current and measuring power continues until either a peak power is found or the current reaches the 20-A maximum.





HOUSEKEEPING TASKS

The remainder of the main loop is devoted to housekeeping tasks, the first of which checks to make sure the load is not dissipating too much power.

Next comes a handler for the front-panel pulse controls, `updatepulse parameters()`, which is (perhaps surprisingly) one of the more complex bits of code in the system. The pulse width and duty cycle are controlled by a rotary encoder hooked to the mbed's internal QEI module, as previously mentioned. The encoder's position is updated once every second by an interrupt-service routine (ISR) that calculates a weighted position based on the encoder's current rotation speed compared to its previous rotation speed. The raw change-of-position data is used as is, is squared, or is cubed to generate the weighted change. This creates the desired speed-sensitive user interface.

Two more key processes, `syncon()` and `syncoff()`, are operating in the background. These use the calculated values of pulse width and duty cycle to generate sync pulses and change the DAC settings. They call each other in ping-pong style, using mbed "Timeout" calls.

As they operate in an interrupt context, I wanted to minimize the time spent in each. Therefore, the DAC settings are not written within these routines (which would have entailed a 600-µs blockage). Rather, the routines simply toggle the LDAC bit to latch pipelined data into the DACs. New data is actually written to the DACs as one of the housekeeping chores in `updatepulseparameters()`.

The next housekeeping module controls the PWM signal that drives the heatsink fans. The heatsink temperature (obtained by the `adc inputs()` module) is mapped to a nonlinear fan-speed value that provides a gentle breeze if the heatsink temperature is in the low 30°C range, rising to full blast as the temperature approaches 50°C. The `fanhandler()` state machine ensures that the fans start up with a brief full-throttle burst to confirm that they don't stall at low speeds.

If the Ethernet module was loaded at initialization, it is necessary to call `Net::poll()` in order for the network state machines to function. After this, the final housekeeping task, `LEDhander()`, mops up the miscellaneous real-time tasks such as updating the LED displays. This routine also counts the number of passes that occur between power-line interrupts: if more than the specified number occur, the power is deemed to have failed and a `powerfail()` routine is invoked. Similarly, the operation of the capture timer is monitored. If it does not interrupt frequently enough (as is the case with zero voltage at the input, for example), then the `LEDhander()` raises a flag for use by the voltage measurement process elsewhere.

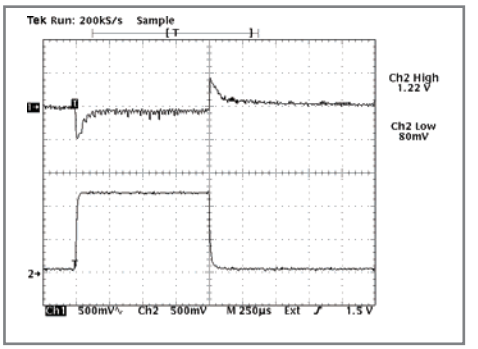


Photo 4—The dynamic response of a power supply is being tested. The scope is triggered by a sync signal as the (VI)sualizer generates a load-current pulse going from 800 mA to 12.2 A (lower trace, scale factor = 10 A/V). The top trace shows the disturbance in the power supply output.

ADDITIONAL CODE

The user interface for the (VI)sualizer comprises five webpages, totaling about 40 KB of HTML and JavaScript code. In addition, there is one external JavaScript file with 9 KB of remote procedure call (RPC) routines and there are two large image files (for the logo and the block diagram) that are fetched from a Picasa account. The webpages use RPC calls to invoke functions on the mbed, which then return data by means of a JSON-compatible string. Code for these pages is available on the *Circuit Cellar* FTP site.

Two new drivers were developed in the course of this project. The first was a driver for the on-chip QEI. I adapted the Cortex Microcontroller Software Interface Standard (CMSIS) driver from ARM for use with the mbed.[3] This

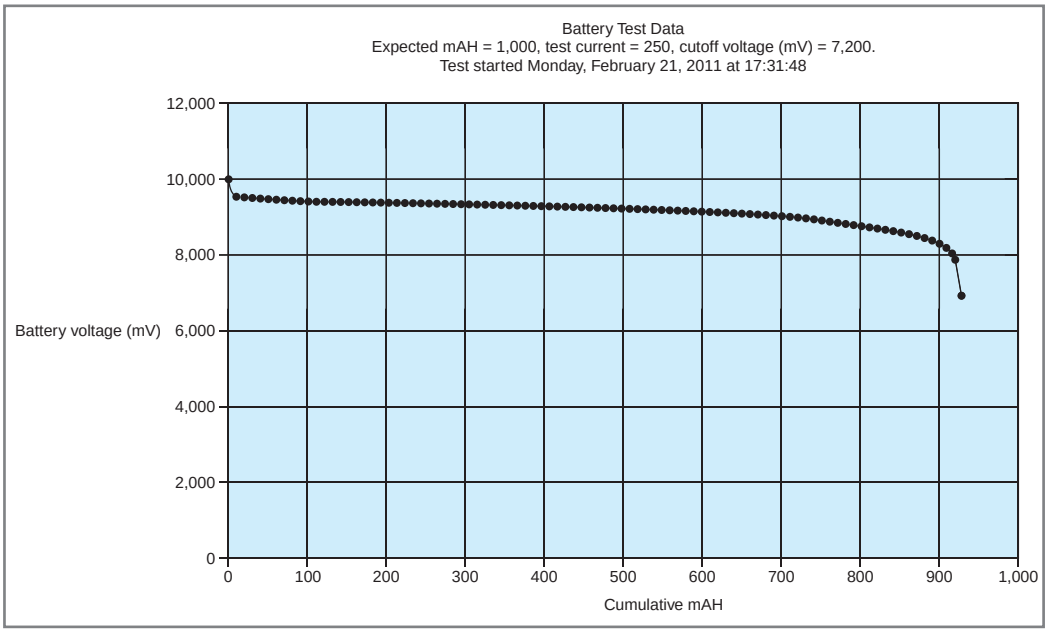


Figure 5—A NiCd battery is being tested to determine its storage capacity. The data was collected by the (VI)sualizer, stored in the flash drive, then later plotted with a spreadsheet program.

entailed enhancements for nested vector interrupt controller (NVIC) support and for user-attached ISRs. One method was added to report direction and another method was added to set the reload time in microseconds.

The second new driver was for the on-chip capture function that is part of the timer module. The driver enables users to specify a count of samples that should be averaged, and then uses an ISR to perform the desired averaging. I took advantage of a hardware quirk to do the summing portion of the averaging with no CPU overhead. Since the driver constrains the sample count to be a power of two, the division part of the averaging is simply a right shift by the appropriate number of places. The overall result is a powerful data-averaging timer-capture function with very low CPU overhead.

RESULTS

Photo 4 shows a scope trace of pulse operation. For this test, a PC-switching power supply was hooked to the (VI)sualizer with a 0.1-Ω current-sampling resistor to ground. The instrument was set to generate pulses starting at about 800 mA and rising to about 12.2 A, with pulse width about 1.0 ms and duty cycle around 1%. We can see that this switcher is well behaved, although it exhibits a visible switching noise when heavily loaded.

Figure 5 shows the voltage versus time data for a nickel-cadmium (NiCd) battery that is discharged at a steady current. This data was captured in the Voltage mode and then charted in Excel.



Photo 5—The results of the automated load finder. The top box shows voltage and current values reflecting the start of a new test. The middle box enables the user to start and stop the test sequence. The bottom box shows the results of the last completed test.

A sample webpage showing the output from a test run in Power mode is shown in Photo 5. I simulated a solar array



Module	Driver			Comments
	Handbook	Cookbook	New	
Ethernet	X			mbed LINK and SPEED signals used to control rear-panel LEDs on the magjack
TCP/IP		X		Network connectivity
HTTP server		X		Webpage delivery
RPC function		X		Webpage user interface tool
Supertweet		X		For IP discovery assistance
ntp client		X		Automated clock setting
USB	X			
MSC file system		X		Data file storage on 4-GB flash drive
mbed Flash				
Local file system	X			Stores webpages
Serial	X			Console interface
PWM outputs	X			Three outputs used (fan speed, two meter drive signals)
Analog inputs	X		X	Used mbed library to initialize, used my own code to make nonblocking reads with integrated averaging and deglitching (all six inputs used).
I²C	X			EEPROM for parameter save/restore
SPI	X			Drive an optically isolated DAC, drive a digital port expander
QEI			X	Developed a full C++ driver based on the CMSIS driver. Supports hardware capture of velocity, user interrupts, etc. Used for pulse/duty cycle control inputs.
Timer block				
Ticker	X			Three instances
Timeout	X			Two instances
Timer	X			Three instances
Capture			X	Developed code to support event capture on Timer 2. Interrupt driven, supports integrated averaging to reduce jitter.
Digital inputs	X			Four signals
Digital outputs	X			Five signals
Interrupt inputs	X			AC line signal for power fail detection

Table 1—Library usage for the project. Software re-use of The mbed Handbook and The mbed Cookbook libraries was key to the rapid development of this project.

by using a 20-V bench power supply in series with a 4-Ω, 50-W power resistor. The (VI)sualizer determined that maximum power was developed at approximately 2.5 A and just below 10 V from the load. And, as expected, it calculated that the optimal load is about 4 Ω. (The displayed figure is about 0.2 Ω low, which probably represents wire losses in the test setup.)

A POWERFUL SOLUTION

I knew from the outset that the (VI)sualizer would require some serious power from its embedded processor. It seemed that the mbed would be a good fit, and indeed that proved true.

What I had not fully appreciated was the degree to which development would be eased by the availability of drivers and libraries from the mbed community. Virtually all of the needed libraries were available from "The mbed Handbook" and "The mbed Cookbook" collections (see Table 1).[4] I can truly say that bringing this design to the same level of complexity would not have been possible in the same time frame had I used a bare-metal processor without the full range of software provided to mbed users.

I see only one drawback. It is so tempting to whip up prototypes with the mbed, that it will be difficult to leave it in situ. I am going to have to tightly screw the cover onto the (VI)sualizer to keep things intact!

Curt Terwilliger won First Prize in the 2011 NXP mbed Design Challenge. He has worked in design and management at several Silicon Valley technology companies. Curt holds two patents, and earned Electrical Engineering degrees from M.I.T. and Stanford. He may be reached via e-mail at his screen name address, hexleyball@gmail.com.

PROJECT FILES

To download the code, go to ftp://ftp.circuitcellar.com/pub/Circuit_Cellar/2011/257.

REFERENCES

[1] S. Morris, "High Temperature Linear Operation of Paralleled Power MOSFETs," Paper presented at HiTEN, the International Conference on High Temperature Electronics Network, 2007.

[2] mbed, The mbed Schematic, 2010, http://mbed.org/media/uploads/chris/mbed-005.1.pdf.

[3] onARM, "The Cortex Microcontroller Software Interface Standard (CMSIS), Version 2.0," 2010, www.onarm.com/cmsis.

[4] "The mbed Handbook," http://mbed.org/handbook/Homepage and "The mbed Cookbook," http://mbed.org/cookbook/Homepage.

SOURCES

MCT6 Optocoupler
Fairchild Semiconductor | www.fairchildsemi.com

HexFET IRF530N
International Rectifier | www.irf.com

MAX232 Line driver/receiver
Maxim Integrated Products | www.maxim-ic.com

MCP4812 Dual-voltage output DAC and MCP6144 Quad op-amp
Microchip Technology, Inc. | www.microchip.com

LM331 Precision voltage-to-frequency converter and LM35 precision centigrade temperature sensor
National Semiconductor | www.national.com

LPC1768 mbed module
NXP Semiconductors | www.nxp.com

NEED-TO-KNOW INFO

Knowledge is power. In the computer applications industry, informed engineers and programmers don't just survive, they *thrive* and *excel*. For more need-to-know information about some of the topics covered in this article, the *Circuit Cellar* editorial staff recommends the following content:

Solar-Powering the Circuit Cellar
Part 1: Preparing the Site
by Steve Ciarcia
***Circuit Cellar* 209, 2007**

Have you considered installing a photovoltaic (PV) system at your home or office building? Steve recently took the plunge. In this series of articles, he describes the entire process, from site preparation to completion. Topics: Solar Power, Photovoltaic, PV, Panel, Inverter, Installation, Pole-Mount

Solar-Powering the Circuit Cellar
Part 2: From the Ground Up
by Steve Ciarcia
***Circuit Cellar* 210, 2008**

In Part 1, Steve described the process of installing his solar PV system. In this article, he covers the massive task of preparing two 11′ schedule 80 poles (8″ in diameter) to support two of the three solar arrays. Topics: Intelligent Energy Solution, Solar, PV, Pole Mount, Wind Load, Concrete, Footing

Go to *Circuit Cellar*'s webshop to find these articles and more: www.cc-webshop.com

AP CIRCUITS
PCB Fabrication Since 1984

As low as...
$9.95
each!

Two Boards
Two Layers
Two Masks
One Legend

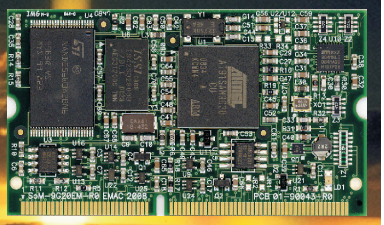
Unmasked boards ship next day!

www.apcircuits.com

VISA MasterCard PayPal IPC MEMBER ASSOCIATION CONNECTING ELECTRONICS INDUSTRIES BBB ACCREDITED BUSINESS

System on Module
Internet Appliance Engine

- Atmel ARM9 400Mhz CPU
- 6 Serial Ports & 2 SPIs
- Up to 60 Digital GPIOs
- 10/100 BaseT Ethernet
- SODIMM Bus Expansion
- SSC/I2S Audio Interface
- Linux with Eclipse IDE
- Real Time Clock Calendar
- SD/MMC Flash Card Interface
- Up to 512 MB Flash & 64 MB RAM
- 2 USB 2.0 Host Ports & 1 Device Port
- 4 10-Bit A/Ds & 6 16-Bit Timer/Counters



2.6 KERNEL
Microsoft Windows CE

The SoM-9G20 uses the same small SODIMM form-factor utilized by other EMAC SoM modules, and is the ideal processor engine for your next design. All of the ARM9 processor core is included on this tiny board including: Flash, Memory, Serial Ports, Ethernet, I2S Audio Interface, PWMs, Timer/Counters, A/D, digital I/O lines, Clock/Calendar, and more. Like other modules in EMAC's SoM product line, the SoM-9G20 is designed to plug into a custom or off-the-shelf Carrier board containing all the connectors and any additional I/O components that may be required. The SoM approach provides the flexibility of a fully customized product at a greatly reduced cost. Single unit pricing starts at $150.

http://www.emacinc.com/som/som9G20.htm

Since 1985 OVER 24 YEARS OF SINGLE BOARD SOLUTIONS

EMAC, inc.

Phone: (618) 529-4525 • Fax: (618) 457-0110 • Web: www.emacinc.com

by Michael Bohn, Tristan Money, Richard Neden, and Keita Sasaki (Canada)

Sleep-Stage Alarm Clock

An MCU-Based System for Processing EEG Data

This microcontroller-based design functions as a sleep-stage alarm clock and an EEG recording system. It uses a neural-network sleep-stage classification system, a data logger, a user interface, and a graphical LCD.

Having problems getting up in the morning? The Sheep Counter design will process your electroencephalo-graphy (EEG) data and wake you up when you are in a light sleep, helping you feel more alert when you arise. This sleep-stage alarm clock uses a neural-network sleep-stage classification system, a data logger to record raw brainwave data to an SD card, a user interface for menu navigation, and a graphical LCD to display sleep information.

Our team (comprising four Camosun College students) was fuelled by a curiosity about the human brain, the most complex "computer" that exists. We began by logging brainwaves and trying to detect sleep stages in real time with a digital signal processing (DSP) algorithm. Despite a lack of formal training in this area, we managed to create a working prototype of our system (see Photo 1) and learned some interesting information along the way.

The design is a sleep-stage alarm clock, automatic sleep-stage classification system, and EEG recording system. We used OpenEEG's ModularEEG for the data-acquisition system. Our sleep-stage classification algorithm was implemented on a Microchip Technology dsPIC33FJ128GP802 microcontroller interfaced to a Ramtron International FM24V10 Serial FRAM. The real-time clock calendar (RTCC), menu system, LCD firmware, and SD card data logging capabilities were implemented on a Microchip Technology PIC18F87K22 microcontroller.

EEG, REM, & R&K

EEG uses small sensors called electrodes to measure the difference in electrical potential between points on the scalp. Neurons produce minuscule voltages when they fire. These voltages can be measured by the sensors to produce EEG signals. Our project uses EEG technology to track a user's sleep patterns.

There are two main categories of sleep: rapid eye movement (REM) and non-REM. Most vividly recalled dreams occur during REM. Non-REM sleep breaks down into four stages: N1, N2, N3, and N4. N1 is a transitional stage between awake and asleep. N2 is associated with light sleep. N3 and N4 are associated with deep, slow-wave sleep. Each sleep stage is defined by its distinct amplitude and frequency characteristics.

Sleep data is usually classified in a clinical environment according the Rechtschaffen and Kales (R&K) rules. According to these rules, sleep EEG is visually classified based on the frequency and amplitude content of a 30-s epoch of sleep EEG.

ALGORITHM DESIGN

We downloaded professionally classified sleep EEG from the PhysioNet database. The format for sleep EEG at PhysioNet is the European Data Format (EDF). Each EEG recording was sampled at 100 Hz. The EDF browser we used was Informer Technologies's EDF checker/viewer Polyman. Each EDF file contains physiological signals for one subject for one full night. There are two EEG signals taken from two

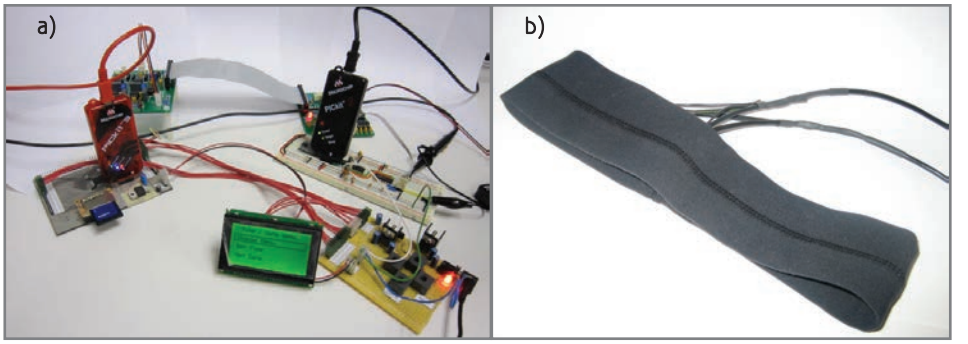


Photo 1—The sleep-stage alarm clock (a) and the EEG headband (b)

K-complexes and Delta	0.415 Hz	to	1.538 Hz
Delta	1.563 Hz	to	3.198 Hz
Theta	3.223 Hz	to	8.594 Hz
Alpha	8.618 Hz	to	10.986 Hz
Spindle	11.011 Hz	to	15.576 Hz
Beta 1	15.601 Hz	to	21.997 Hz
Beta 2	22.021 Hz	to	37.476 Hz

Table 1—Sleep stages

electrode pairs within each EDF file. We chose the signal taken from F_{PZ} and C_Z.

We used a macro program called Auto Hot Key (AHK) to automate the process of separating each classified 30-s epoch of sleep EEG from each full night's sleep EEG recording. We saved each of these 30-s epochs in separate files. Each 30-s epoch was sorted according to its respective classification.

Our data set needed to be composed of a sufficient number of samples to build a robust algorithm. Based on an article by E. Oropesa et al from the International Computer Science Institute, we decided that 400 samples from each sleep type would be adequate.[1] We individually separated and sorted approximately 5,600 30-s epochs of sleep EEG.

We converted the sorted EDF data into a format that could be opened in GNU Octave using the wfdb2mat program provided in the PhysioToolkit software package from PhysioNet. We used an AHK script to automate this process.

The base and gain information is contained in the EDF file header for each EEG recording. This information was required to convert the 16-bit signed integer sample values to their corresponding physical units (µV). We loaded each 30-s epoch of sorted EEG data into GNU Octave and scaled it according to its appropriate base and gain values.

Next, we performed a 4,096-point Fast Fourier Transform (FFT) on each 30-s epoch of sleep EEG using the GNU Octave FFT function. Using the FFT, we determined the power spectral density (PSD). To present the data to the artificial neural network (ANN), each 30-s epoch of sleep EEG had to be reduced to 13 features. (For more information, refer to the sidebar, "Artificial Neural Networks.") We determined the total power within specific frequency bands, which formed the first seven features for the epoch.

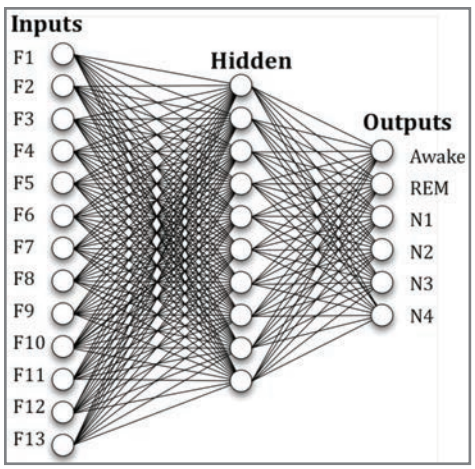


Figure 1—Neural net layout

We followed the method used by Oropesa et al. These frequency bands correlate to the R&K rules for the manual scoring of sleep EEG. We summed the PSD indices corresponding to the frequency bands listed in Table 1.

We calculated an additional six features based on the sum of the power within all seven previously defined features and ratios of these previously defined features. We determined these features for each of the 5,600 sorted 30-s epochs of sleep EEG.

We applied a logarithmic transform to each of the classified data sets containing the reduced 30-s epochs of sleep EEG to obtain a normal distribution of the data. These 13 features make up the 13 input vectors for the ANN. We incorporated a corresponding target vector associated with each input vector. The target vector consisted of five zeros and a one in a position according to the specific sleep stage.

We used Cranfield University's AMOR Neural Networks package for the CRAN R language and environment for statistical computing. We followed the neural network structure developed by Oropesa: an input layer comprised of 13 neurons, a hidden layer comprised of 10 neurons, and an output layer comprised of six neurons (see Figure 1). The hidden layer is fully connected to the input layer and the output layer is fully connected to the hidden layer. We used a hyperbolic tangent sigmoid activation function for

ARTIFICIAL NEURAL NETWORKS

Artificial neural networks (ANNs) are computational tools that model the biological operational characteristics of the human brain. ANNs are used to solve problems involving pattern recognition. A diverse range of problems across a host of disciplines can be solved using the robust computational capabilities of ANNs.

In general, an ANN consists of an interconnected network of nodes or neurons. There are a wide variety of customizable, problem-specific network configurations available. Each connection or synapse between neurons has a multiplicative weight value associated with it. Each neuron has a transfer function or activation function and bias value associated with it. The input neurons are assigned values that are specific characteristics or features derived from the raw input data. Each output neuron corresponds to a specific stage or correct classification of the input.

In our case, during the training phase, the ANN is presented with data that has been classified by a human. The weight and bias values associated with each neuron and connection are initially random values. An algorithm adjusts the weight and bias values associated with each connection and neuron until the output of the network correlates to the correct classification of the input data as closely as possible. A large volume of classified input data and a multitude of training iterations are usually required before a network can learn to accurately recognize specific patterns.

Listing 1—Neural network code segment

```
for(neuron_index=0; neuron_index < 10; neuron_index++)
}
     // increment through weights
     for(weight_index=0; weight_index < 13; weight_index++)
     {
          a += hidden_neurons[neuron_index].weights[weight_index] * transform[weight_index];
     }

     a += hidden_neurons[neuron_index].bias;

     hidden_neurons[neuron_index].v0 = a_tansig * tanh(a * b_tansig);

     a = 0;
}
```

the hidden neurons and a linear activation function for the output neurons. We used a gradient decent back-propagation training method with momentum.

We trained and simulated the ANN using the AMOR package ANN functions. Using 1,000 randomly selected samples from our data set of 2,400 samples, we trained the network. We used the remaining 1,400 samples to test our result. We picked the best result from multiple runs to implement in our sleep-stage recognition algorithm. Our best result achieved approximately 84% overall agreement. We anticipate that this result could be improved with careful optimization.

We saved the neural network's weight and bias values associated with our best training run. We used these values in the subsequent C implementations of the algorithm.

ALGORITHM AS SOFTWARE

We implemented the entire preprocessing procedure using the C programming language and verified our values against the results we obtained using GNU Octave. The dsPIC33FJ128GP802 did not have hardware to do floating-point math. To address this issue and to make our C code as portable to the microcontroller as possible, we used an FFT that used fixed-point math. There are no complex data types in C. Therefore, we implemented the FFT by separately calculating the real and imaginary components of the complex numbers.

We used a MathWorks fixed-point FFT tutorial to tailor a 4,096-point fixed-point FFT for our application. We precalculated the twiddle factors to improve efficiency. The twiddle factors are the complex points between zero and pi on the unit circle. The twiddle factors are multiplicative constants used in the butterfly calculation, which is the principal computational element of the FFT. The twiddle factors are stored in memory and are referenced by their array indices in the butterfly calculation. We precalculated and stored the bit-reversed address order in memory. The bit-reversed address order is used in the time-domain decomposition stage of the FFT.

Using GNU Octave, we designed the algorithm and scaled the signed integer sample values to the physical units (µV) prior to computing the FFT. We performed division to scale the signed integer sample values. When we divided the signed integer sample values, the values were converted to floats. The GNU Octave FFT uses floating-point math. However, microcontroller hardware is not available to do floating-point math. Consequently, we used a fixed-point FFT on the microcontroller to improve efficiency. Therefore, we determined a scaling factor for the base and gain settings for each night's worth of sleep EEG after the FFT and PSD were calculated.

By comparing the PSD results from our program with the PSD results from GNU Octave, we determined the scaling factor and achieved an appropriately scaled PSD using fixed-point calculation. One disadvantage of using fixed-point arithmetic is that errors are introduced due to truncation. However, these errors had a marginal impact on the accuracy of the PSD.

We determined the 13 features in our software C implementation using the same method described above. We took the weight and bias values associated with each neuron in the ANN from the network created using the AMOR Package for CRAN R and hard coded them into an array of structures. We referenced the source code from the AMOR Neural Networks package to tailor the ANN implementation to suit our application.

The calculations in the implementation of the ANN were simple. For each hidden neuron, we multiplied each input with the weight corresponding to the input's connection with the hidden neuron. We did this for each input connection coupled to that hidden neuron. We added the products of the inputs and weights associated with each hidden neuron and stored the sum in variable a. The bias value for each hidden neuron was then added to the total stored in variable a. The calculation used in the activation function is described below. The result of the calculation is stored in variable v0. We repeated this process for each of the output neurons, multiplying each v0 value associated with each hidden neuron by the weight corresponding to its connection with the output neuron. The activation function associated with each output neuron is linear as opposed to hyperbolic tangent sigmoid for the hidden neurons (see Listing 1).

ALGORITHM AS FIRMWARE

Microchip Technology has some efficient FFT functions that take advantage of the dsPIC's built-in bit-reversed addressing. However, since there are only 16 KB of internal RAM on the dsPIC, these functions limited us to a 512-point FFT. We were forced to interface to external RAM to fulfill the memory requirements for a 4,096-point FFT. We used the 1-Mb Ramtron FM24V10 serial FRAM to satisfy our memory


Where the next big thing
meets the light of day.
Somewhere a back room conversation has caused a burst of inspiration.
And it will be unveiled here. Come see a bright idea light the way to tomorrow.
2012 International CES
THE GLOBAL STAGE FOR INNOVATION
PRODUCED BY CEA
Tuesday, January 10–Friday, January 13, 2012
Las Vegas, Nevada | CESweb.org
REGISTER NOW

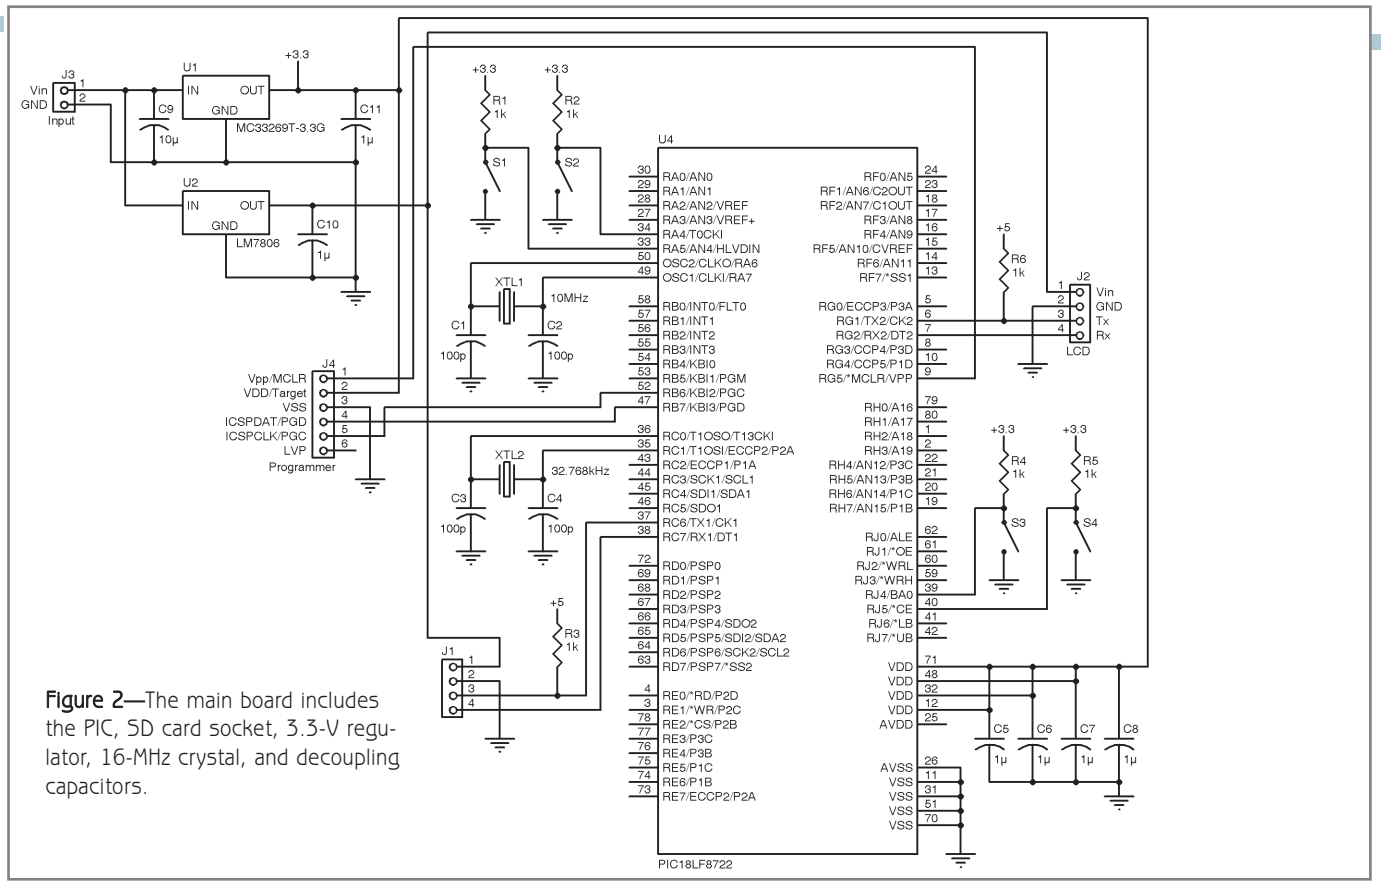


Figure 2—The main board includes the PIC, SD card socket, 3.3-V regulator, 16-MHz crystal, and decoupling capacitors.

requirements.

We serially transmitted the real and imaginary components of the twiddle factors as well as the bit-reversed address order to the dsPIC where they were stored in external RAM.

In the C implementation of the FFT function described in the previous section, we referenced the real and imaginary components of the input sample values and twiddle factors based on the array index of the required value. Since we were using external RAM, we needed to reference the input sample values and twiddle factors based on their address locations in memory instead of their array index. We therefore modified the software C implementation to incorporate referencing the input values and twiddle factors based on their addresses in RAM as opposed to their array indices.

The dsPIC only performs multiplications on integers up to 32 bits. The twiddle factors are 16-bit signed values. The product of the twiddle factors multiplied by the inputs would often exceed the range of a signed data type (-2^{15} to $2^{15}-1$). We resolved this issue by multiplying each negative signed integer by –1 to make it positive. We cast the signed integer as an unsigned integer. The range of an unsigned integer is -2^{16} to $2^{16}-1$. This is sufficient to accommodate the product of the two values. The appropriate sign was reintroduced after the bit-shift-right operation was performed as part of the fixed-point calculation. The code for the feature extraction and ANN was fully portable from the software C implementation of the algorithm to the microcontroller.

Implementing an FFT as we have done on the microcontroller is particularly inefficient. Transmission of the 30-s epoch of sleep EEG to the microcontroller at 115,200 bps until the classification is complete takes approximately 6 s. However, when we apply the classification algorithm on

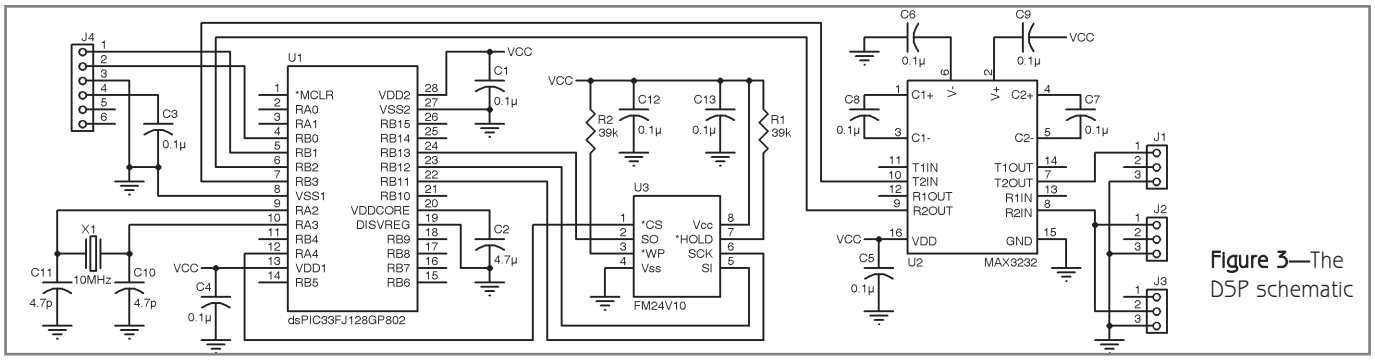


Figure 3—The DSP schematic

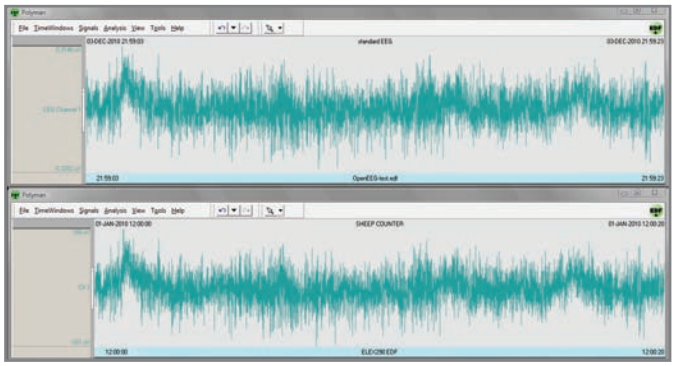


Photo 2—EDF storing EEG data. The top shows the EEG recording using a PC-based software product. The bottom shows the EEG recording using the Microchip Technology PIC18F87K22 microcontroller.

real-time EEG from the Open-EEG, we should easily have the capacity to do the signal processing between receive interrupts within the 30-s time frame before the next 30-s epoch arrives. If efficiency becomes pertinent, an ARM architecture or an FPGA would be an ideal solution.

MAIN BOARD

A Microchip Technology PIC18F87K22 microcontroller is used as the clock, data logger, and user-interface processor. We chose this microcontroller for its high clock speed (64 MHz) and large amount of program space (128 kB). The 64-pin thin quad flat pack (TQFP) PIC18F85K22 microcontroller would have been more suitable (less pins), but it was not available from Microchip at the time of development. We required a breakout board to access pins on the PIC's TQFP80 surface-mount package. The breakout board includes the PIC, SD card socket, 3.3-V regulator, 16-MHz crystal, and decoupling capacitors. Figure 2 shows the mainboard schematic and Figure 3 shows the DSP schematic. We created a daughterboard to house the power input, a 9-V regulator for the LCD, a 5-V regulator for level conversion, a 5-V to 3.3-V level converter, four buttons, and connections to the breakout board and LCD module. The two boards and LCD are connected to each other with crimpless header cables.

FILE SYSTEM

We used the Microchip Technology Memory Disk Drive File System (MDDFS) for FAT16/32 file writing. This is a convenient library for writing and reading secure digital, compact flash, or USB flash memories. We chose FAT because it is widely accepted as the standard file system for memory devices, making it compatible with most computers. The MDDFS library takes up about 80 kB of the 128 kB of program memory on the PIC, which may exceed the program memory capacity of other PIC models. The size of the MDDFS can be reduced if features such as FAT32 support, directories, or file searching are disabled. We left these options enabled since they may be used in a later version of the Sheep Counter.

DATA, DISPLAY, MENU, & RTCC

We used the EDF for storing EEG data. This is a widely used format supported by many software packages such as



Polyman and Brainbay (see Photo 2). The EDF header file stores information about the user, the data-acquisition system, and the data-logging equipment used to create the recording. The header is relatively short (512 bytes) because the EEG uses only one channel. The header size can quickly increase if additional EEG channels are used.

We used a 128 × 64 graphical LCD with a serial backpack for the main display. The serial backpack reduces the amount of interfacing pins required from 18 to two and provides built-in drawing functions and fonts. The initial firmware on the serial backpack was slow for writing characters. Therefore, we used the improved firmware from Geeknet's SourceForge. We further modified the firmware to ignore null characters, have a larger input buffer, and correct some font errors.

Four debounced buttons provide control over a multilevel menu system. The four menu options: Time and Date Display, EEG Graphing and Data Logging, Set Backlight Brightness, and Set Time and Date. The Set Time and Date menu goes into another menu level to provide separate time and date setting functionality.

We used the PIC18F87K22's built-in RTCC. We clocked the RTCC with a 32.768-kHz external crystal oscillator. We initialized the RTCC by setting the RTCEN register. The clock has registers for the year, month, day, and time. The user can set the clock using four navigation buttons.

GRAPHING EEG WAVEFORM

The microcontroller serially transmits sleep EEG sample data to the 8-bit PIC (see Photo 3). The 8-bit PIC stores the sleep EEG samples in an array. We write the average of every eight samples of sleep EEG data to the LCD using the built-in line-drawing functions for the LCD. We transmit a start character to the LCD followed by the line-draw function command. The x and y coordinates of each sample correspond to the line endpoints. We write the average of every eight samples of sleep EEG to the LCD as a y coordinate. We then increment the x coordinate corresponding to each y coordinate by one. The LCD functions attach all the averaged sleep EEG sample points with a line to create an EEG waveform. The EEG waveform was then scaled to ensure a proper fit on the LCD. Refer to Figure 2 to review the circuitry.

MOVING FORWARD

We are going to continue to refine our project into a more seamlessly integrated system. If efficiency becomes pertinent, we might use a 32-bit microcontroller or an FPGA. We will optimize the neural network to improve our results and improve the accuracy of the fixed-point FFT on the microcontroller.

Because the original research conducted during this project was successful in producing the desired results, we would like to make a contribution to academic understanding in this field. In addition, we would like to acknowledge the work of Edgar Oropesa, Hans L. Cycon, and Marc Jobert for the field of sleep-stage classification.

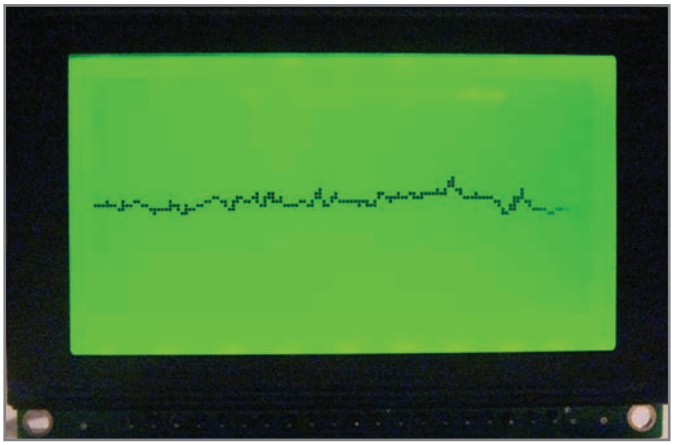


Photo 3—A sleep EEG graph on a serial LCD

Authors' note: We graduated from the Camosun College Electronics and Computer Technology program in December 2010. We would like to thank the Camosun College Electronics Engineering Department for their support and assistance: Mel Dundas, Wayne Mayes, Joyce van de Vegte, and Amar Bouallouche. We would also like to express our gratitude to Moshi Charnell in the Math Department and Kate Wellburn in the English Department for their help with the project.

Michael Bohn (mike.bohn@gmail.com) lives in Victoria, BC, and is currently working for Calgary-based Comtech Communications Technologies. Michael is interested in EEG, biofeedback, and entrainment. He also produces and composes music.

Tristan Money (tristanm@shaw.ca) splits his time between Port Alberni and Vancouver, BC. He works in automation/IT for automationX Industrial Solutions, Inc. Tristan is looking forward to working with industrial control systems and finishing his electric motorcycle.

Richard Neden (lord_richard_the_great@hotmail.com) lives in Victoria, BC, and is interested in renewable energy and biomedical engineering. Richard is apprenticing to be an alarm technician and maintains interests in biomedical engineering and renewable energy.

Keita Sasaki (sakeincanada@hotmail.com) lives in Victoria, BC, and is currently working as sensor assembler at AML Oceanographic. Keita enjoys wireless communication systems and digital signal processing.

REFERENCES

[1] E. Oropesa, H. Cycon, and M. Jobert, "Sleep Stage Classification Using Wavelet Transform and Neural Network," International Computer Science Institute Technical Report TR-99-008, 1999, ftp://ftp.icsi.berkeley.edu/pub/techreports/1999/tr-99-008.pdf.

RESOURCES

Cranfield University, Applied Mathematics & Operational Research (AMOR), "Neural Networks," 2006, www.amorg.co.uk/module_NN.html.

Geeknet, Inc., SourceForge, http://sourceforge.net/projects/serialglcd.

GNU Operating System, GNU Octave, www.gnu.org/software/octave.

MathWorks, www.mathworks.com.

OpenEEG, The ModularEEG, http://openeeg.sourceforge.net/doc/modeeg/modeeg.html.

PhysioNet, physionet.org.

SOURCES

EDF Checker/view Polyman
Informer Technologies | www.informer.com

dsPIC33FJ128GP802, PIC18F87K22, and PIC18F85K22 Microcontrollers, TQFP80 breakout board and Memory Disk Drive File System (MMDDFS)
Microchip Technology, Inc. | www.microchip.com

FM24V10 serial FRAM
Ramtron International Corp. | www.ramtron.com

NEED-TO-KNOW INFO

Knowledge is power. In the computer applications industry, informed engineers and programmers don't just survive, they *thrive* and *excel*. For more need-to-know information about some of the topics covered in this article, the *Circuit Cellar* editorial staff recommends the following content:

—

DIY Signal Generation
Create a Signal Source for a Variety of RF Applications
by Neal Martini
***Circuit Cellar* 219, 2008**

Neal describes how he built an inexpensive and accurate variable-frequency signal generator on a two-sided printed circuit board. The system delivers frequencies from 0 to 3 GHz with a frequency accuracy of less than 200 Hz. Topics: Signal, Generation, RF, Variable Frequency, DDS, PLL, VCO, LCD, Keypad, Phase, dsPIC30F3010A, MAX5026, AD9851, ADF4113, Board LayoutZX05-43LH, V600ME10-LF

—

Go to *Circuit Cellar*'s webshop to find this article and more: www.cc-webshop.com





by Robert Lacoste (France)



CE Marking

A Process to Ensure Product Conformity

What is CE marking and why is it important? CE stands for *Conformité Européenne*, which is French for "European Conformity." CE marking states that the manufacturer of the product ensures that the product is compliant with the essential requirements of the relevant European directives. Learning about CE marking is useful for engineers who are both using and designing products.

My column usually deals with deep technical subjects. This month, I'll make an exception: I'm going to talk about regulations and standards. I admit it may not be the most fun topic on earth, but I still encourage you to read this article. I'll do my best to avoid long lists of legal-oriented stuff and stay as engineer-oriented as usual.

Why an article on regulation and, more specifically, on European CE marking? Because, there are lot of misunderstandings and exaggerated fears or unconscious optimism around this subject. If you don't live in Europe, or if you have never designed a product to be sold on this continent, I think you'll find this article a useful introduction to the topic.

One caution: A short article like this one could only be a simplistic presentation of tens of thousands pages of regulations, standards, and procedures. Moreover, I am neither a lawyer nor an expert on standards, so don't take my word as the truth. And, always double check with official sources. Anyway, I have reasonable experience in that field, so a significant part of what I will present will likely be true, even if standards change every day.

CE MARKING DEFINED

Let's begin. What is CE marking? CE means *Conformité Européenne*, which is French for "European Conformity." Physically, it is a CE logo, at least 5 mm tall, that must be visible on nearly all products sold in Europe (see Photo 1). In fact, only some product classes are in the scope of this regulation; but as this is the case for all electrical and electronic equipment, radio devices, and machines, I guess you should be concerned.

As CE marking is mandatory, it is absolutely not a kind of "quality mark." All electronic devices must have it, even if they are poorly designed. It is also not a certificate of origin. Products sold in Europe must be CE marked, even if they were designed in the U.S. and manufactured somewhere else, such as in an eastern country. CE marking is also not a proof of independent assessment. As I will explain, a manufacturer is usually not obliged to ask a third party to do the tests.

CE marking is a conformity marking. It states that the manufacturer of the product ensures the product is compliant with the essential requirements of the relevant European directives, meaning the directives applicable to that product. Just read that sentence again. CE marking is a declaration of the manufacturer. It engages its responsibility. Period. So, this is very different from Underwriters

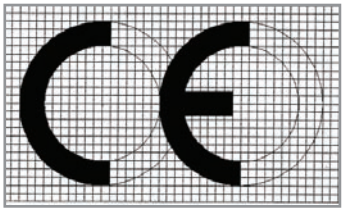


Photo 1—The CE logo, currently found on all products sold in Europe

Laboratories (UL) or Canadian Standards Association (CSA) approvals, which imply the verification of the product's characteristics by an independent testing house. Here, the manufacturer declares that it has done the required conformity assessment, possibly on its own, and that the product is compliant.

I will present what an essential requirement is in a minute, but first I will clarify some definitions. CE marking is required for all electronic products "sold" in Europe. More exactly, the legislation says, "placed on the market or put into service" rather than "sold." This means that it is not applicable to an experimental prototype staying on your bench as it is not yet on the market. But, it is theoretically required, for example, if you send some of your products for free to early adopters or if you rent your products.

The term "manufacturer" also needs some explanation. An importer or a distributor can also be considered a manufacturer. This is the case, in particular, when products are delivered under the distributor's name or when a product is not CE marked by the manufacturer. The importer or distributor will then assume the legal responsibility for the CE marking process.

ESSENTIAL REQUIREMENTS?

So, what are these "essential requirements" CE marking must conform with? Basically, these are generic requirements that the product must comply with. They are defined in so-called "European directives," which can be considered European laws. In addition to generic paperwork-oriented directives, three of them are applicable to a large number of electrical devices. They deal respectively with electrical safety, electromagnetic compatibility (EMC), and radio equipment and telecommunications terminal equipment (R&TTE). Moreover, there are directives on hazardous-substances restrictions, waste management, and energy efficiency. However, there are also plenty of European directives each applicable to certain classes of products (e.g., machineries, toys, medical devices, lifts, etc.). Table 1 gives you an overview.

Let's consider an example of an essential requirement. The directive 2004/108/EC on EMC says: "An equipment shall be so designed and manufactured, having regard to the state of the art, as to ensure that: (a) the electromagnetic disturbance generated does not exceed the level above which radio and telecommunications equipment or other equipment cannot operate as intended; (b) it has a level of immunity to the electromagnetic disturbance to be expected in its intended use which allows it to operate without unacceptable degradation of its intended use."

Similarly the so-called "low voltage" 2006/95/EC safety directive says: "The Member States shall take all appropriate measures to ensure that electrical equipment may be placed on the market only if, having been constructed in accordance with good engineering practice in safety matters in force in the Community, it does not endanger the safety of persons, domestic animals or property when properly installed and maintained and used in applications for which it was made."

However, theses directives don't say at all *how* to achieve compliance with these requirements. This is the responsibility of the manufacturer! So, at its roots, CE marking doesn't imply compliance with any precise measurement values or performances. Once again, it says that the manufacturer asserts that the product is compliant with these generic requirements, that's all. By the way, all CE European directives are brief and available for free download. If you are interested, feel free to read them, but you will usually be disappointed by the level of detail provided!

HARMONIZED STANDARDS?

Fortunately, the manufacturer usually does not have to deal directly with the European directives. There are European harmonized standards. These documents, developed by technical organizations—such as Comité Européen de Normalisation (CEN), which is the European Committee for Standardization; Comité Européen de Normalisation Électrotechnique (CENELEC), which is the European Committee for Electrotechnical Standardization; or the European Telecommunications Standards Institute (ETSI), which is a standards organization for the telecommunications industry—are actual detailed engineering-oriented standards that provide information a designer is looking for (e.g., test plans, thresholds, etc.). The rule is as follows: if there are harmonized stan-

CE directives applicable to a large number of electronic products	
2006/95/EC	Low voltage electrical equipment (safety)
2004/108/EC	Electromagnetic compatibility (EMC)
1999/5/EC	Radio equipment and telecommunications terminal equipment (R&TTE)
CE directives specific to some product classes	
2009/48/EC	Toys
2006/42/EC	Machinery
93/42/EEC	Medical devices
204/22/EC	Measuring instruments
94/9/EC	Equipment used in potentially explosive atmospheres (Atex)
95/16/EC	Lifts
Etc.	
Other directives applicable to nearly all products, not directly linked to CE marking	
2002/95/EC	Restriction of use of hazardous substances (RoHS) in electrical and electronic equipment
2002/96/EC	Waste from electrical and electronic equipment (WEEE)
2009/125/EC	Ecodesign of energy-related products (EuP)

Table 1—There are now only 20 or so European directives linked to CE marking, each applicable to certain product classes. This table also lists the three directives linked to environment protection.

EN 41003:1998	Particular safety requirements for equipment to be connected to telecommunication networks
EN 60065:2002	Audio, video, and similar electronic apparatus; Safety requirements
EN 60204-1:2006	Safety of machinery, electrical equipment of machines; Part 1: General requirements
EN 60215:1989	Safety requirements for radio transmitting equipment
EN 60335-2-5:2003	Household and similar electrical appliances; Safety, Part 2–5: Particular requirements for dishwashers
EN 60335-2-59:2003	Household and similar electrical appliances; Safety, Part 2–59: Particular requirements for insect killers
EN 60598-2-7:1989	Luminaires; Part 2: Particular requirements, Section 7: Portable luminaires for garden use
EN 60730-2-14:1997	Automatic electrical controls for household and similar use; Part 2–14: Particular requirements for electric actuators
EN 60825-1:2007	Safety of laser products; Part 1: Equipment classification and requirements
EN 60950-1:2006	Information technology equipment; Safety, Part 1: General requirements
EN 60950-21:2003	Information technology equipment; Safety, Part 21: Remote power feeding
EN 61010-1:2001	Safety requirements for electrical equipment for measurement, control, and laboratory use; Part 1: General requirements
EN 61558-1:2005	Safety of power transformers, power supplies, reactors, and similar products; Part 1: General requirements and tests
Etc.	

Table 2—Some examples of the harmonized standards that can be used for the Low Voltage directive 2006/95/EC. From the hundreds available, the difficulty is finding out which standards are applicable to your specific product.

dards applicable to your product—say, for EMC—and if your product is compliant with these standards, then your product presumably complies with the associated essential requirements. If you reread that sentence, you will understand that there is no formal obligation to comply with any standard for CE marking. However, if you don't use standards, your life will be far more difficult. You will have to prove that you comply with very generic assumptions. For example, for EMC, you will need to prove that your product has no risk of interference with any other product. Good luck!

It is the manufacturer's responsibility to find out which standards are applicable, and there are plenty of them. The European Commission keeps a database of all harmonized standards associated with each directive, but the lists are long. For example, there are around 300 harmonized standards associated with the generic "low-voltage" safety directive. See Table 2 for some examples. Selecting which standards are applicable to your product is really a job for an expert. My advice—even if it is not mandatory—is to consult an experienced testing house, or, if possible, an official notified body (which are testing houses accredited to do conformity assessments for a given list of standards) even just to list the applicable standards. This should be done as early as possible in the development cycle and will make your life far more comfortable.

What are directives and standards actually about? I will, of course, not be able to present here specific directives dealing with lifts or medical devices, but I will try to give you an overview. As listed in Table 1, the three directives that are applicable to nearly all electronic devices deal with: safety (Low Voltage directive), electromagnetic compatibility (EMC directive), and radio and telecommunications (R&TTE directive). Let's take a look at each directive.

LOW VOLTAGE DIRECTIVE

The so-called Low Voltage directive (2006/95/EC) covers all electrical equipment with input or output voltages between 50 and 1,000 VAC or between 75 and 1,500 VDC. Therefore, it is applicable, in particular, to all 110/230-V line-powered equipment, but not to small battery-operated systems. Be careful, as there is a specific case: it is also applicable to all radio transmitters. So, your last ZigBee-enabled sensor is concerned, even if its power source is a 1.5-V AAA battery.

Let's consider a computer peripheral—say, a printer. If you look at the list of harmonized standards associated with the Low Voltage directive, you will find that the EN60950-1 standard is applicable to "information technology equipment." Seems a good match. You will then have to buy a copy of this EN60950-1 standard (directives are free of charge, but standards usually aren't) and read it carefully. In this standard, you will find plenty of precise requirements on electrical insulation and safety (insulation distances, safety earth, double versus single insulation, surge suppressors, etc.). The safety isn't restricted to the electrical side. For example, you will also have to ensure that the product is mechanically safe (stable, no harmful parts, etc.) as well as thermally safe (no accessible hot spots, resistant to flame hazard, etc.).

If you have a worldwide market, you are lucky. This standard is, in fact, aligned with the international IEC60950-1 standard, as well as the UL60950-1 standard, so you will only have to do the job once. This is fortunately often the case. Mutual recognition between standard agencies are better and better. You will have to check the details.

EMC DIRECTIVE

The other big part of CE marking is the EMC directive (2004/108/EC). As for safety, there are hundreds of applicable standards, depending on your product. Keeping the same example of a computer peripheral, the best match is usually the EN55022 standard—"Information Technology Equipment, Radio Disturbance Characteristics, Limits and Methods of Measurement"—associated with the more generic EN61000-X-X EMC immunity standards. The EMC standards cover five different topics. Let's consider each one.

Radiated emissions: The product is switched on in an anechoic chamber, in front of a receiving RF antenna. It shouldn't radiate anything above the standard-defined thresholds from tens of megahertz up to gigahertz.

Conducted emissions: Similarly, if the product has any external wiring (e.g., power cord, I/O cables, etc.), it shouldn't conduct any noise or spurious signals above given limits through these cables. Usually, the results of this test are linked to the radiated test: if there are conducted spurious signals, then the cable will probably act as an antenna and radiate them, too.

Radiated immunity: This test is the opposite of radiated emissions. Still in the anechoic chamber, the product is exposed to high levels of electromagnetic fields (simulating, for example, a nearby mobile phone), of course with specified levels, frequencies, and modulations. The device must not fail during the test. More exactly, it should behave as stated in its documentation.

Conducted immunity: The same kind of test is also done on external cables, injecting RF noise and spikes on them and checking the behavior of the product.

Electrostatic discharge immunity: Lastly, the product must not be destroyed and/or put in a non-recoverable state (depending on its classification) when exposed to electrostatic (ESD) spikes of some kilovolts. This is usually the last test done on the prototype, especially if it is unique. Something will surely fail if a metallic, accessible part is not properly grounded.

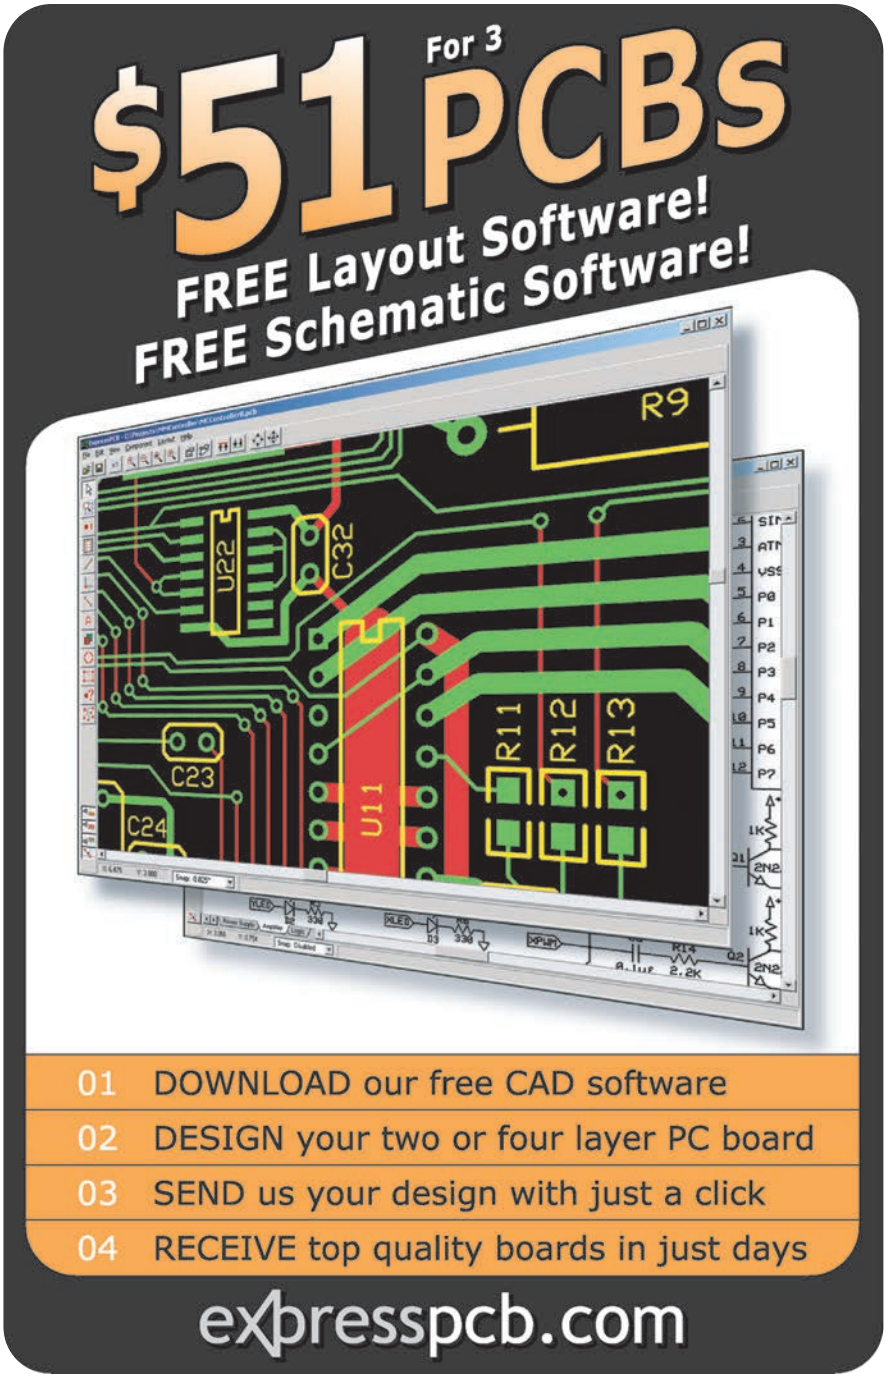
Parts of these tests are similar to some Federal Communications Commission (FCC) standards; but unfortunately, the levels are not identical. A conscientious testing house can do the tests once and consider the product compatible with both standards, but this should be anticipated.

R&TTE DIRECTIVE

The last generic directive applicable to several classes of electronic products is the Radio Equipment & Telecommunications Terminal Equipment (R&TTE) directive (1999/5/EC). This directive is applicable to all radio transmitters, but also to radio receivers and wired telecommunication devices. As in the case of safety or EMC, there are plenty of related harmonized standards, such as the EN300328 standard applicable to wide-band 2.4-GHz devices (e.g., Wi-Fi). These standards force you to design "good" RF devices for an optimal use of the radio bandwidth. A transmitter must transmit at the good frequency and power even if the battery is low. It must not pollute the adjacent channels, and it must not transmit spurious signals elsewhere in the RF spectrum. Lastly, it must usually check if the channel is free before transmitting.

Similarly—even if it is less obvious—a "bad" receiver is not allowed. This will force the use of a higher-power transmitter, so reducing the efficient use of the RF spectrum.

Lastly, devices with a significant transmission power (above some tens of megawatts) will also have to comply with recent electric and magnetic field (EMF) limits related to potentially



harmless RF signals on human tissues. You will usually have to prove that there can't be more than 20 mW of power dissipated in any sample of 10 g of human tissues. These tests, called specific absorption rate (SAR) tests, are unfortunately quite long and quite expensive.

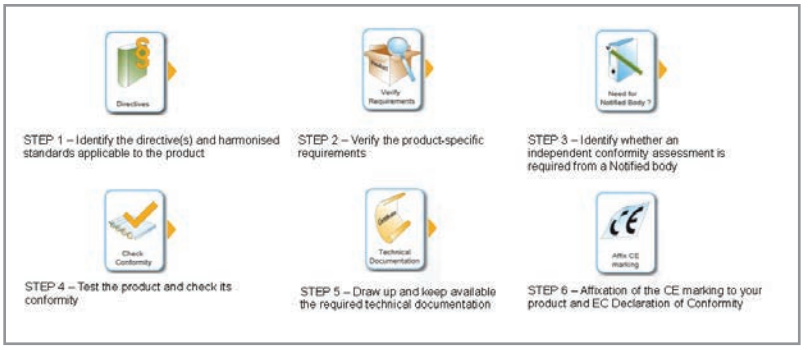


Photo 2—Here are the six steps required by the CE marking procedure

WHAT ABOUT RoHS, WEEE, & EuP?

I'm sure you have heard of environmentally oriented regulations that have emerged in Europe these last few years. Even if these directives are not actually linked to CE marking, they are mandatory, so let's briefly cover the main ones.

First, there is the Restriction of Hazardous Substances (RoHS) directive (2002/95/EC), which covers lead, mercury, cadmium, etc. In a nutshell, only RoHS-compliant products can be sold in Europe, except for very specific types of products (e.g., medical devices, spare parts for older systems, etc.). So, there isn't an associated mandatory marking. Just buy only lead-free solder and RoHS-compliant components.

The second mandatory environmentally linked directive is the Waste from Electrical and Electronic Equipment (WEEE) directive (2002/96/EC). It is associated with a specific marking (a barred garbage bin) and basically states that the manufacturer will manage and pay for the recycling of the product at the end of its lifecycle through defined procedures.

Last but not least, there is the codesign of Energy-Related Products (EuP) directive (2009/125/EC). Its target is to improve the energy efficiency of power-hungry devices as well as to push suppliers to analyze the full lifecycle of their products. For the moment, it is applicable only to selected products (e.g., domestic lightning, TV sets, electric motors, etc.). The list is growing quickly, so stay tuned! For example, all new TV sets introduced in mid-2012 to the European market will have to use less than 1 W in Standby mode and, when in use, less than 16 W + A × 3,4579 W/dm^2, with A the screen size in square decimeters. Do the math on your TV and measure if you are close to that limit. This will show you that there is still some work for electronic engineers!

RESPONSIBILITY & PAPERWORK

Let's briefly consider the administrative side of CE marking. As explained, the manufacturer of a product is responsible for affixing the CE mark on its own product. However, the CE mark states that the manufacturer has taken some mandatory steps prior (see Photo 2). First, as discussed, the manufacturer must determine which directives the product needs to comply with and which standards are applicable to the product. Then, the manufacturer needs to carry out a conformity assessment of the product. Practically speaking, this means that you have to select the applicable standards and execute the corresponding test plans. In rare cases, a notified body must be involved. This is stated in the directives and concerns only "risky" product classes (e.g., pressure-related devices, explosive environments, medical devices, etc.). In the other cases, nothing prevents you from doing the tests in house. However, usually it is more efficient to subcontract the testing to a dedicated test house because you may not have all the required equipment and experience. Also, nothing prevents you from executing only part of the required tests, or even not relying on any applicable standard. But, in that case, you will need to explain (and document) how you can ensure that your product is still 100% compliant with the applicable essential requirements. And, you can be sure that it will be a far more complex task than to simply apply a predefined and standardized test plan.

After the conformity assessment, the manufacturer also needs to set up a technical file. It must include all details of the product design (i.e., schematics, PCB layout, bill of materials, mechanical design, user manuals, test reports). This document, which is the product's bible, must be written in an official European language and must be kept by the manufacturer at least 10 years after the end of the commercialization of the product. This document must be available to authorities in case of controls, and must be physically available somewhere in Europe. For non-European manufacturers, this is usually done through an authorized representative, which could be a distributor or even an attorney if the manufacturer does not have 100% trust in its distributer (as the technical file can enable duplicating the product).

Lastly, the manufacturer must write an EC declaration of conformity, which is a one-page document stating that the product is CE compliant, describing the applicable essential requirements and standards, and providing the manufacturer's name and address. This document must be signed and made available to each customer. Then, you can apply the CE mark on the products and ship them.

Is it finished? No. Products are always alive. Your design will evolve. Some components will be obsolete and will require small redesigns, etc. Through all these incremental changes, the manufacturer must ensure that the CE marking is not impacted. At minimum, you will need to document the changes in the technical file and document why you can ensure that the changes have no impact on the CE conformance. This may not be trivial, in particular when EMC is concerned, so you may need to conduct new full or

partial conformance assessment in order to still commit on the product's conformance. You also need to prove that the shipped products are always in line with the technical file, which will require either manufacturing testing or quality-assurance procedures. Life is interesting, isn't it?

CE MARKING ILLUMINATED

I admit that CE marking is, as all regulations, a complex and often boring subject. There are two ways to see it: you can either focus only on the negative side—more constraints, more costs, more delays, and more paperwork—or on the positive side—thanks to such frameworks, the products we are using are by far safer and more environmentally friendly than those of 20 years ago. Choose your side!

In closing, I must once again emphasize that such a column only gives you a simplified view of the subject. Don't take it for granted, and double check with an expert or your preferred testing house. I hope I have neither simplified my explanations nor provided you with the wrong ideas. My only goal was to give you enough background to consider that CE marking isn't on your dark side anymore.

Robert Lacoste lives near Paris, France. He has 22 years of experience working on embedded systems, analog designs, and wireless telecommunications. He has won prizes in more than 15 international design contests. In 2003, Robert started a consulting company, ALCIOM, to share his passion for innovative mixed-signal designs. You can reach him at rlacoste@alciom.com. Don't forget to write "Darker Side" in the subject line to bypass his spam filters.

RESOURCES

European Commission Enterprise and Industry, "6 Steps to CE Marking for Your Product," http://ec.europa.eu/enterprise/policies/single-market-goods/cemarking/downloads/further_information_en.pdf.

———, Affix CE Marking, http://ec.europa.eu/enterprise/policies/single-market-goods/cemarking/professionals/manufacturers/affix-ce-marking/index_en.htm?filter=11.

———, European Standards, List of References of Harmonised Standards, http://ec.europa.eu/enterprise/policies/european-standards/documents/harmonised-standards-legislation/list-references/#ch2.

———, FAQ, http://ec.europa.eu/enterprise/policies/single-market-goods/cemarking/faq/index_en.htm.

European Council Directives, http://eur-lex.europa.eu.

New Approach Standardisation in the Internal Market, www.newapproach.org.

T. Williams, *EMC for Product Designers*, Newnes, 2007.

Wellkang Tech Consulting, "How Many CE Marking Publications?" www.ce-marking.org/how-many-directives.html.

NEED-TO-KNOW INFO

Knowledge is power. In the computer applications industry, informed engineers and programmers don't just survive, they *thrive* and *excel*. For more need-to-know information about some of the topics covered in this article, the *Circuit Cellar* editorial staff recommends the following content:

—

BITs and BITEs
Tests to Ensure System Safety
by George Novacek
***Circuit Cellar* 251, 2011**

Built-in test equipment is an important component in safety-critical systems. Design engineers must always consider what's needed for safe operation and easy maintenance, and then take measures to adequately balance safety and over-testing. Topics: Watchdog, System Safety, Built-In Test, Equipment, Power Up, Testing, Monitoring, Load

—

SMT Manufacturing
Take a Board from Prototype to Production
by Zack Gainsforth
***Circuit Cellar* 208, 2007**

Zack walks you through the full development cycle of a USB-to-serial adapter for a programmable knob. The design is intended for car computer systems where it can be programmed to support multiple functions. Now you too can take a similar design from prototype to production. Topics: SMT, Development Cycle, Design Capture, Haptic Knob, PR-1000, USB, Adapter, Schematic, Layout, Manufacture, Fabrication, Electronic Design Automation (EDA) Software, Gerber Files, Panelizing, Assembly, Kitting, Test, Rework

—

Mobile Phone Book
M16C/62P-Based Data Backup System
by Carlos Cossio
***Circuit Cellar* 190, 2006**

Carlos's M16C/62P-based system enables you to back up your cellular phone book records on the microcontroller's flash memory, a PC, or a SIM smart card. Topics: SIM Smart Card, GSM, Phone, M30626FHPGP, M16C/62P, APDU, Communications, TPDU

—

Go to *Circuit Cellar*'s webshop to find these articles and more: www.cc-webshop.com

If we haven't convinced you that the Circuit Cellar Archive...

OVER TWO DECADES OF ENGINEERING EXCELLENCE
At Your Fingertips



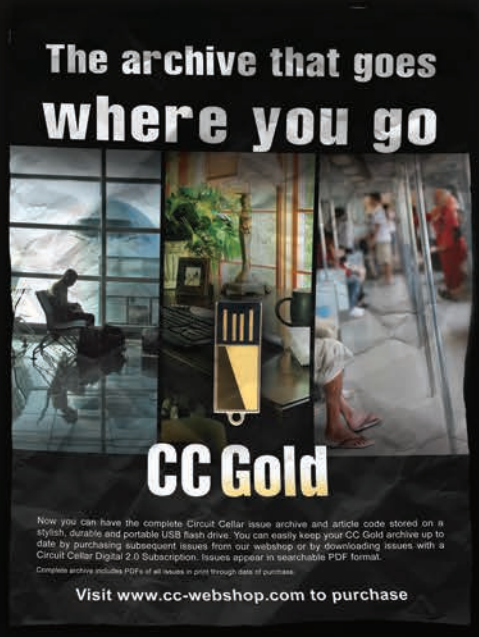
Has over two decades of engineering excellence

Which key opens the most doors?



Now you can have the complete Circuit Cellar issue archive and article code stored on a stylish, durable and portable USB flash drive. You can easily keep your CC Key archive up to date by purchasing subsequent issues from our webshop or by downloading issues with a Circuit Cellar Digital 2.0 Subscription. Issues appear in searchable PDF format.

Opens doors



Goes where you go

Compliments your style





Is the ultimate archive



Then maybe this will help.
Visit www.circuitcellar.com/bigdeal

Now you can have the complete Circuit Cellar issue archive and article code stored on a stylish, durable and portable USB flash drive. You can easily keep your archive up to date by purchasing subsequent issues from our webshop or by downloading issues with a Circuit Cellar Digital 2.0 Subscription. Issues appear in searchable PDF format.

Complete archive includes PDFs of all issues in print through date of purchase.
(For display purposes) USB drive shown is larger than the actual size.

by Bob Japenga (USA)



Getting Started with Embedded Linux (Part 2)

Choosing a Platform for Your System

You know when to choose Linux for an embedded design. This article details the platform types that can run embedded Linux. With this information, you can move ahead and choose the proper foundation for your embedded system.

I remember my first engineering project from my senior year in college. Our task was to design a weather satellite, and my team was responsible for the control system. I remember sitting with my teammates over a set of high-level requirements and a blank sheet of paper. Four months later, our design was complete.

Now, 40 years of projects later, I remember this as one of the few that started with a blank sheet of paper. Since most real projects we work on are built based on some previous work, we rarely have complete control. Almost everything I have done in my career has been accomplished by standing on the shoulders of the giants before me, seeing what others have done, and building upon that. By the way, most of us attribute this metaphor to Isaac Newton. But he was quoting the 12th-century French neo-Platonist philosopher Bernard of Chartre.

What, you may ask, does this have to do with choosing a platform for hosting an embedded Linux system? As we look at the options available, I understand that most of us don't start with a blank sheet of paper. Choices will be made for you by others who will have selected the target cost of the system, the delivery schedule, or the processor architecture. I hope to help you make the choices that you do have control over. I'll start with two options that most of you will never have in order to help you understand the scope of what is available with embedded Linux.

THE MOST COMPLEX OPTION

Imagine you were developing a brand new processor, starting from a completely blank sheet of paper—a new whiz-bang machine, unlike anything else ever developed. Or, perhaps you want Linux to run on an existing processor, one with no Linux support. How would you build a version of Linux for this processor? You could start by visiting the official website for the Linux kernel which is maintained by the Linux Kernel Organization. From there, you could download the latest stable or long-term kernel. At this point, you would need to choose one or two processors that were at least somewhat like yours, look at what they did, and develop all the support code necessary to make your processor run Linux. You would be developing memory-management code, timers, and memory interfaces, all patterned after the Linux kernel design paradigm you see in other processors's Linux code. Eventually, you would also need a toolchain for this new processor, but that is another discussion.

HELPFUL DEFINITIONS

Kernel: the software that is the bridge between hardware and the software applications, providing access to and management of the hardware and software resources

Toolchain: the software that is used to create other software, usually including a text editor, a compiler, a linker, and a debugger

Distribution: a specific package of the Linux operating system that contains the kernel, the toolchain, and the application software

Platform: the hardware for hosting an embedded Linux system

Architecture	Notes
Alpha	You guessed it—Digital Equipment's PC killer
ARM	One of the most popular embedded architectures
CRIS	Code Reduced Instruction Set (CRIS): CPU used in network servers (discontinued in 2011)
FR-V	Fujitsu's FR-V microcontroller family
H8300	Renesas H8/300 Series: high-speed processors featuring powerful bit-manipulation instructions, ideally suited for real-time control applications
I386	X86 family of CPUs
IA64	The Itanium architecture of CPUs, specializing in explicit instruction-level parallelism
M32R	A 32-bit RISC architecture CPU developed by Mitsubishi for embedded systems
M68k	The Motorola 68000 CPU architecture
MIPS	The Microprocessor without Interlocked Pipeline Stages (MIPS) CPU is a reduced instruction set computing (RISC) CPU developed by MIPS Technologies
PPC and PPC64	Power PC 32- and 64-bit CPUs
PA-RISC	Precision-architecture RISC CPU developed by Hewlett Packard
S/390	IBM S/390 series of CPUs
Super H and Super H64	A 32/64-bit RISC architecture of microcontrollers for embedded systems developed by Hitachi
SPARC and SPARC64	The Scalable Processor ARChitecture (SPARC), developed by Sun
V850	Renesas's CPU architecture, targeted at real-time processing applications
X86-64	X86 64-bit family of CPUs

Table 1—Some of the processor families/architectures supported by Linux

THE NEXT OPTION

Just as I have rarely started with a blank sheet of paper, you will probably not be asked to put Linux on a brand-new processor family. You may, however, be called upon to put Linux on a processor that was built by the giants before you. In that case, you will just add a new processor within an existing family of processors. Table 1 provides a list of some the families currently supported by the Linux Kernel Organization. This list will give you an idea of the breadth of support Linux has in the processor community. Under each of these families, there can be scores of individual processors.

If you put Linux on a new processor, your software will be added to the list of supported processors for your family (called an architecture). As before, you would be modifying all the key hardware interfaces, modifying the toolchain, and then building the necessary Linux applications.

GETTING TO REAL EXAMPLES

I would hazard a guess that few of you fall into either of these first two categories. More likely, you are putting Linux on a system with a processor that's already supported.

If your processor is already supported, use the Linux configuration tool provided for the kernel and select your processor, the various peripherals (such as the PHY), and all of the glue logic for your project. If these are supported, then build the kernel and the associated support applications (i.e., file system, network options, etc.), and away you go.

Hold on. What if your peripherals and glue logic are not supported? Great question. At this point, you are no longer standing on the shoulders of giants. Instead, your hardware designer is standing on your feet. If the hardware is already cast in copper, you are stuck and will need to write your own drivers, expending a lot of time and money. To simplify your task, all the peripherals and glue logic you place around the processor should preferably be selected from the list of those supported. This means you should try to be involved in the hardware design early on. Fortunately, a lot of devices are supported. Unfortunately, I have never found a definitive, up-to-date list, because it is constantly changing. My approach has been to take the Linux kernel source code I have downloaded and search the driver tree for the particular analog-to-digital converter, real-time clock, PHY, or I/O expander I am using and try to find out if it's supported. Just use the first few characters of the device part number and see what you find. If the device is supported, you'll need to find out whether that driver is supported by your specific processor. For this, you can use the Linux configuration tool. Sometimes, the chip vendor will say that the device is supported under Linux. It may even provide the driver. For example, Analog Devices recognized the importance of providing Linux drivers, and has provided a comprehensive list of its peripherals with Linux drivers. I expect others will soon follow.

Thus far, I've glossed over another important part of your project. All the supporting application software needs to be compiled for your particular processor. Things like the C library (assuming you are programming in C), the file system, ntp, telnet, or a web server will all need to be built for your particular processor. I've also skipped over another important issue: your toolchain. Stay tuned for more on that in my next column.

OpenEmbedded OPTION

The Linux Kernel Organization approach is not for the faint of heart. OpenEmbedded is slightly better. This open-source project was created to enable embedded-system developers to create a complete Linux distribution. There is an active community of developers, just like you and me, who are building embedded systems. Unlike the Linux Kernel Organization's kernel.org, which is used by the entire Linux community, OpenEmbedded is specifically targeted to create the kind of systems you and I work with.

OpenEmbedded supports a wide variety of architectures and processors, but there are some unsupported architectures, such as the DEC Alpha. You can choose your processor following a set of recipes and build not only the kernel, but the toolchain and the supporting Linux applications.

We have used OpenEmbedded to move a project built around a reference design (see below) to a later version of the kernel. At the end of the project, the entire package was well configured and the process was repeatable. (It reached Level 2 on the Capability Maturity Model continuum of software development, which was developed by the Software Engineering Institute at Carnegie Mellon University to help improve software reliability.)

The disadvantage of using OpenEmbedded is that there is a steep learning curve to learn how to use its tools. It can take weeks to get Linux up on your board.

REFERENCE DESIGN OPTION

If you're starting to feel dizzy, just hang in there. It's getting easier, not harder. Another option for a hardware and system designer who hopes to use Linux is to choose a reference design, which is supplied by chip or board manufacturers. My company has used reference designs from Cogent Computer Systems, Texas Instruments, and Atmel. They provide schematics, application notes, a bill of materials, and a complete Linux distribution. This approach is quick and easy. You can begin writing software on a development board, which is usually similar to your own board at the computer architecture level. You can design your board with the exact same or similar parts to those used in the reference design, and then use the same Linux distribution you used on the development board. We have had Linux running on our proprietary board within one day when we use this approach.

A disadvantage to this approach can occur if the manufacturer does not support the processor or the supporting chips. If you look at the Analog Devices webpage I referenced earlier, you'll note that the company has indicated whether or not each device is supported. For products with long life spans and many software iterations, this is critical.

Another problem can happen when you want to add a device that is only supported in a newer kernel. For example, we wanted to add a new Wi-Fi USB device and found that it required a newer kernel or back porting the driver—both painful processes. If all of your other drivers are not supported in that newer kernel (e.g., because the chip manufacturer never put them into the standard kernel), adding that device will be painful. When this happened to us, we were severely limited in our ability to add new devices the customers demanded.

SINGLE-BOARD COMPUTER OPTION

If your budget can tolerate it, perhaps the simplest platform to build a Linux project on is a single-board computer (SBC) or a system on a module (SOM) that comes with Linux already on it. There are literally thousands of such SBCs and SOMs. With this option, you can forget about spending any extra time getting embedded Linux to work, and benefit immediately from its power. In other words, you can get straight to what makes your company unique: your application.

Obviously, this comes at a higher recurring cost. However, I have found that it works well in two situations: when the production volume shipped is low (less than 1,000 per year) or when the time to market is critical. In one of our projects, a company needed to hit the ground running. To make money in its business model, it needed to sell more than 3,000 per year, but it was willing to take a loss on the first few thousand. So, we designed the system around an SBC and delivered the prototype within nine weeks, and a fully qualified version—one that met the standards of Underwriters Laboratories (UL), the Federal Communications Commission (FCC), and the American National Standards Institute (ANSI)—within six months. Then, after it established market presence and its volumes skyrocketed, we redesigned the company's system around an ARM9 reference design from Atmel. All the application code now runs on both systems, despite the differences between the two processors.

Finally, another one of our customers built their system around an SBC using an X86 architecture using a compact flash for memory. In this case, you can take a standard Linux distribution—such as Ubuntu, Fedora, or Red Hat—drop it on the card, tweak the basic input/output system (BIOS) to boot from the card, and go. This customer knew nothing about Linux, but developed a successful embedded Linux application by booting from flash memory this way without our help!

ENDLESS OPTIONS

Price, schedule, and history can all affect the platform options available to you. But there are many options. There is much we haven't yet covered, and what we have covered has only been in thin slices. But then, that is what this column is all about. See you next time, when we'll talk about toolchains and licensing issues.

Bob Japenga has been designing embedded systems since 1973. In 1988, along with his best friend, he started MicroTools, which specializes in creating a variety of real-time embedded systems. With a combined embedded systems experience base of more than 200 years, they love to tackle impossible problems together. Bob has been awarded 11 patents in many areas of embedded systems and motion control. You can reach him at rjapenga@microtoolsinc.com.

RESOURCES

Analog Devices Wiki, Linux Driver, 2011, http://wiki.analog.com/resources/tools-software/linux-drivers-all?force_rev=1.

The Linux Kernel Organization, www.kernel.org.

OpenEmbedded, 2011, www.openembedded.org.

SOURCE

ARM9 Reference design
Atmel Corp. | www.atmel.com

ICbank
Gateway to Electronics

MP3P DIY KIT, Do it yourself

(Include Firmware Full source Code, Schematic)

• myPIC



Only
$160 qty 100
$220 qty 1

Powerful feature

- MP3 Encoding, Real time decoding (320Kbps)
- Free charge MPLAB C-Compiler student-edition apply
- Spectrum Analyzer
- Application: Focusing for evaluation based on PIC
- Offer full source code, schematic

Specification
Microchip dsPIC33FJ256GP710 / 16-bit, 40MIPs DSC
VLSI Solution VS1033 MP3 CODEC
NXP UDA1330 Stereo Audio DAC
Texas Instrument TPA6110A2 Headphone Amp(150mW)
320x240 TFT LCD
Touch screen
SD/SDHC/MMC Card
External extension port (UART, SPI, I2C, I2S)

• myWave (MP3 DIY KIT SD card Interface)



Only
$150 qty 100
$200 qty 1

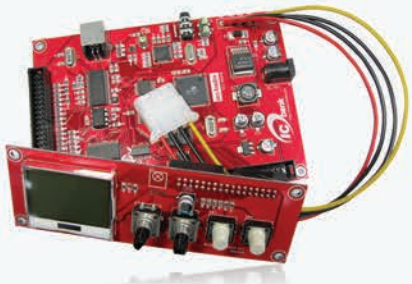
Powerful feature

- Play, MP3 Information, Reward, forward, Vol+/-
- Focusing for MP3 Player
- SD Card interface
- Power: battery
- offer full source code, schematic

Item	Specification
MCU	Atmel ATmega128L
MP3 Decoder	VS1002 / VS1003(WMA)
IDE Interface	Standard IDE type HDD(2.5", 3.5")
Power	12V, 1.5A
LCD	128 x 64 Graphic LCD
Etc	Firmware download/update with AVR ISP connector

• myAudio (MP3 DIY KIT IDE)

Only
$180 qty 100
$220 qty 1



Powerful feature

- Play, MP3 Information, Reward, forward, Vol+/-
- Focusing for full MP3 Player (Without case)
- IDE Interface
- Power: Adapter
- Offer full source code, schematic

ICbank
Gateway to Electronics

www.icbank.com

by Ed Nisley (USA)



Stepper Motor Drive (Part 3)

Torques

This article describes a simple dynamometer that can reveal how a stepper motor performs while turning different mechanical loads. While it's not completely accurate, the dynamometer should give you an appreciation of the real-world performance of motors.

Now that you have seen how a microstepping driver regulates the current in a stepper motor's windings, it's time to examine the end product: the torque that turns the shaft.

In this column, I'll describe a simple dynamometer that can reveal how a stepper motor performs while turning different mechanical loads. Although it doesn't produce laboratory-grade accuracy, the dynamometer should give you an appreciation of the real-world performance of motors.

DYNAMOMETER MECHANICS

The notion of using a dynamometer to measure engine power dates back to 1821, when Gaspard de Prony invented his eponymous de Prony Brake: a spring balance attached to an arm on a friction brake around the motor shaft. The spring balance indicates the braking force and, multiplied by the arm length, the motor torque. Multiplying torque and speed gives the motor's power: power = torque × speed.

Tightening a friction brake around a motor shaft works much better on engines producing output power measured in kilowatts or horsepower than on the small stepper motors I'm using. The force applied by such a brake will vary due to the inevitable irregularities in the braking surface, and I expect that those variations would be relatively large compared to a small stepper's average torque.

Because stepper motor ratings give the motor's torque, rather than its output power, rearranging the equation suggests a different approach. The ratio of two more easily measured values gives the torque:

$$\text{torque} = \frac{\text{power}}{\text{speed}}$$

Speed measurement requires nothing more than an optical interrupter similar to the one I used in my October column, "Stepper Motor Drive (Part 2): Microstepping," and converting mechanical energy into electrical energy makes power easy to measure. A spreadsheet can handle the math: all we need are data points.

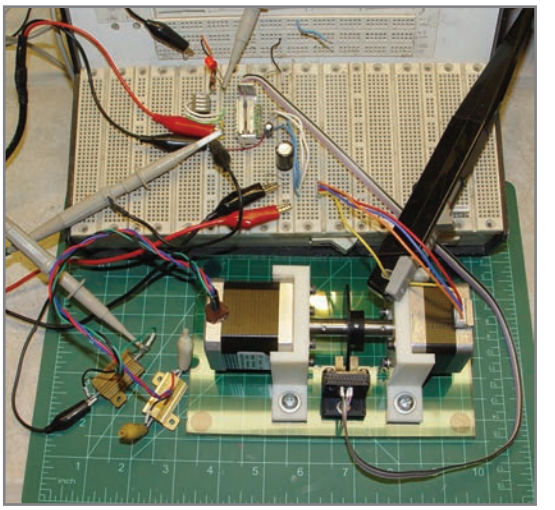


Photo 1—A tiny Pololu A4988 Stepper Driver Carrier board under the heatsink turns the smaller Minebea motor, with a Tektronix current probe monitoring its winding current. The larger ACT motor acts as a generator dissipating power in the two aluminum power resistors. The sync wheel and optical interrupter on the coupling shaft produce a 1/rev pulse for oscilloscope triggering.

With that in mind, the dynamometer in Photo 1 consists of two NEMA 17 stepper motors connected by a shaft. The smaller stepper, the same house-labeled Minebea 17PM-J034 you saw my October column, is the "motor under test." The larger stepper, a Changzhou ACT Motor Company ACT 17HS5425, acts as a generator that dumps electrical power into the two aluminum-body power resistors to its left. I have a large assortment of such resistors, making it straightforward to vary the power dissipated at a given voltage.

The solderless breadboard holds a Pololu A4988 Stepper Motor Driver Carrier under the small aluminum heatsink, with DIP switches setting the microstepping mode. Figure 1 from my October column shows the wiring details, with the note concerning the unsuitability of breadboards for high-power wiring remaining in effect.

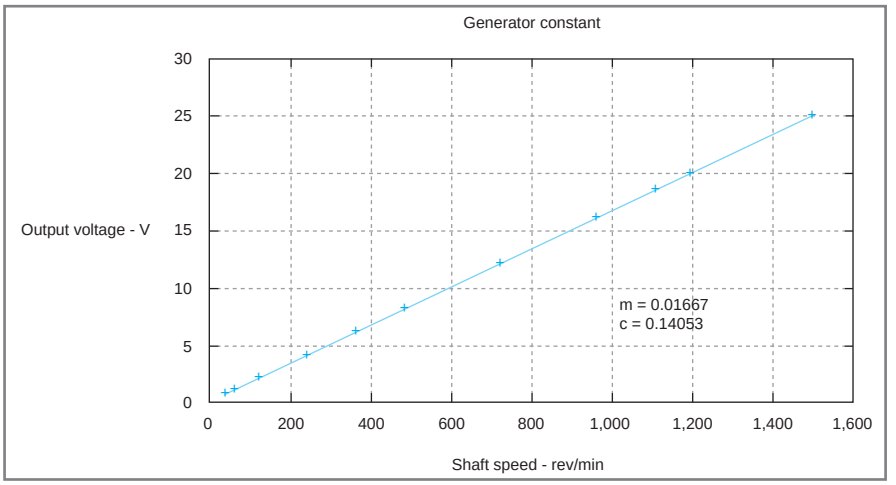


Figure 1—The slope of the voltage versus speed line equals the larger stepper motor's open-circuit generator constant Kg. The non-zero intercept value shows that my measurements aren't quite perfect: the motor does not have a 140-mV output at 0 rev/min. With no load on the output, the smaller stepper motor stalls at about 1,500 rev/min.

The two alligator clips just below the DIP switch connect a pulse generator to the Step signal input. The pulse frequency thus determines the motor speed and I can vary that speed very slowly to eliminate the influence of mechanical acceleration on the measured values.

The notched disk on the connecting shaft interrupts the IR beam in the optical switch near the bottom of the picture so that a pullup resistor on the breadboard produces one pulse for each shaft revolution. Because we're concerned only with complete motor revolutions, there's no need for the Arduino sync generator you saw last time.

Measuring power in watts (the unit recalls James Watt, who improved the earliest steam engine to the point of usability) comes naturally to engineers, but we think of motor speed in revolutions per minute (rev/min or rpm) or per second (rev/s), rather than the SI unit of radians per second (rad/s). Converting rev/min to rad/s gives the motor torque in newton meters (N · m):

$$N \cdot m = \frac{W}{rpm \times 2\pi/(60\ s/min)}$$

Figure 1 plots the 17HS5425 motor-as-generator's open-circuit peak output voltage against the shaft speed. The slope of that amazingly straight line gives the *generator constant* K_g as 16.7 mV_{pk}/rpm. The value of K_g depends on the generator's construction, so you must measure it for any other motor you'll use as a generator. For example, I measured the Minebea J034 motor's K_g at 14 mV_{pk}/rpm.

Strictly speaking, the 17HS5425 motor is an alternator because it produces an AC output voltage. You could use an internally commutated DC motor to produce pulsating DC, but I used two stepper motors for mounting simplicity and the ease of interchanging the roles of motor and generator.

Although the output voltage varies linearly with the shaft speed, the power dissipated in the load resistor varies as the square of the speed:

$$W = \frac{E^2}{R} = \frac{(K_g \times rpm)^2}{R}$$

At very low speeds, the output voltage approaches zero, which means the generator's winding resistance limits the maximum available power dissipation. Fortunately, stepper motors produce reasonably constant torque at low speeds, so extrapolating the results from the values measured at moderate speeds produces useful results.

Operating the dynamometer requires some hands-on skill. Photo 2 shows the voltage across a 4-Ω resistor as I slowly increased the step frequency until the motor stalled. The output voltage reads about 3.25 V_{pk} just before the J034 motor cannot drive the increasing load, the rotor stalls, and the voltage drops. It doesn't drop to zero, however, because the stepper driver continues to produce winding currents

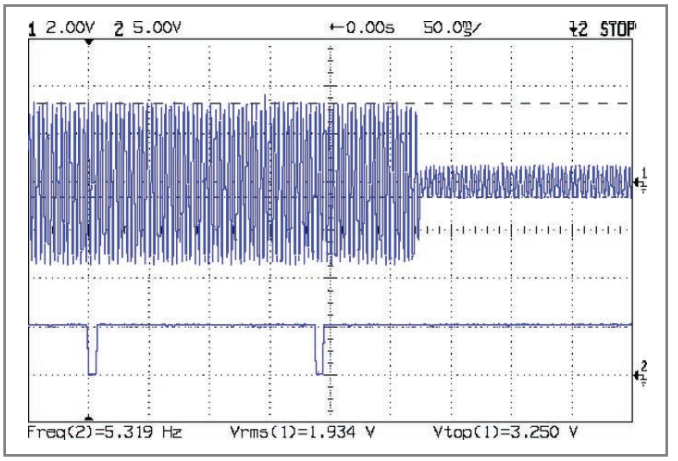


Photo 2—The motor stalls when the load imposed by the generator exceeds the torque produced by the stepper driver. The rapid voltage drop and missing 1/rev sync pulse shows that the rotor stops almost immediately after the motor stalls, leaving a small voltage due to the vibrating motor shaft.

that vibrate the shaft around its stalled position.

I calculated the power dissipation using the root-mean-square voltage, rather than the peak voltage, to find the DC equivalent power and eliminate the effect of small waveform distortions. The oscilloscope calculates the rms value directly from the displayed waveform, which made that an easy choice. The scope processes the entire trace, however, so I took measurements when the motor stalled "off screen" where it wouldn't affect the result. The manually calculated value for the 3.25 V_{pk} shown in Photo 2 is 2.30 V_{rms}.

The current through the 4 Ω resistor is:

$$575\ \text{mA} = \frac{2.30\ \text{V}}{4\ \Omega}$$

The 17HS5425 winding resistance adds 1.1 Ω to the 4.0 Ω external resistor for a total circuit resistance of 5.1 Ω. The total power dissipation for one winding is: 1.7 W = $(0.575\ \text{A})^2 \times 5.1\ \Omega$.

Because the two windings have voltages and currents in quadrature, the torque depends on the peak current in a single winding, which is √2 larger than the rms value. The corresponding peak power is a factor of two larger than the value calculated from the rms voltage or current.

Pop quiz: show why a √2 increase in current through a resistor boosts the power dissipation by a factor of two.

The 1/rev pulses in the lower trace of Photo 2 correspond to a rotational speed of 5.32 rev/s = 319 rev/min = 33.4 rad/s, which is a full-step frequency of 1,060 step/s. The motor torque is therefore:

$$100\ \text{mN} \cdot \text{m} = \frac{2 \times 1.7\ \text{W}}{\frac{33.4\ \text{rad}}{\text{s}}}$$

Although the Minebea catalog doesn't list the 17PM-J034 motor, a presumably similar 17PM-J049 motor shows about 190 mN·m of pull-out torque at 300 rpm and its rated 0.7-A drive.

However, the 6.2-Ω winding resistance of the J049 motor dissipates 3 W at its rated current. The 2.1-Ω windings in my J034 suggest a rated current of 1.2 A would run at the same power dissipation. Scaling the 1 A produced by the Pololu driver on my breadboard to 1.2 A would produce about 120 mN·m of torque, although you'll see shortly that such extrapolations may be misleading.

The Minebea data also assumes a 24-V power supply, rather than the 12 V in my test setup. As you saw in my August and October columns, a higher drive voltage produces faster current changes in the windings, with a corresponding increase in torque output. I also expect the Minebea engineers would optimize the test conditions to wring the best possible performance from their motors.

Given that I don't know the expected torque, I think the results are in the right ballpark. Regardless of any absolute errors, this dynamometer can produce meaningful relative data for a variety of drive and load conditions.

Homework: build a friction-based de Prony Brake and measure the torque for a motor with known characteristics. Bonus: estimate the range of possible errors.

TORQUE VS. SPEED

An ideal stepper motor driver would impose constant currents in the windings during each microstep, regardless of the rotational speed and, thus, the microstep's duration. That means the motor provides constant power during the microstep and, because torque is inversely related to speed at a given power, as the speed increases the available torque must decrease.

Secondary effects, including the driver's inability to produce a rectangular current step due to winding inductance, cause the torque to drop off more rapidly as the microstep rate increases. Beyond a few thousand rpm, the motor produces a tiny fraction of its rated torque.

Conversely, at speeds up to a few hundred rpm where the driver can impose a nearly rectangular current step, the motor produces almost constant torque that is limited by the armature's ability to convert electrical current into magnetic field, regardless of the speed. Stepper motors thus tend to be useful in relatively high-torque, low-speed applications with shaft speeds ranging from near zero to a few hundred rpm.

As described above, the dynamometer data gives the motor's torque at its stalling speed under load. Figure 1 shows that the generator voltage depends linearly on the speed, so the resistance must vary to impose a specific load at a given speed. That would require matched power rheostats, which I don't have, so I chose matched pairs of power resistors from my heap and raised the speed until the generator's increasing voltage dissipated enough power to stall the motor. That means the data points will be irregularly spaced due to the torque variations.

Figure 2 plots the maximum torque available at each speed in 1/8 microstepping mode with the driver applying 1-A_{pk}

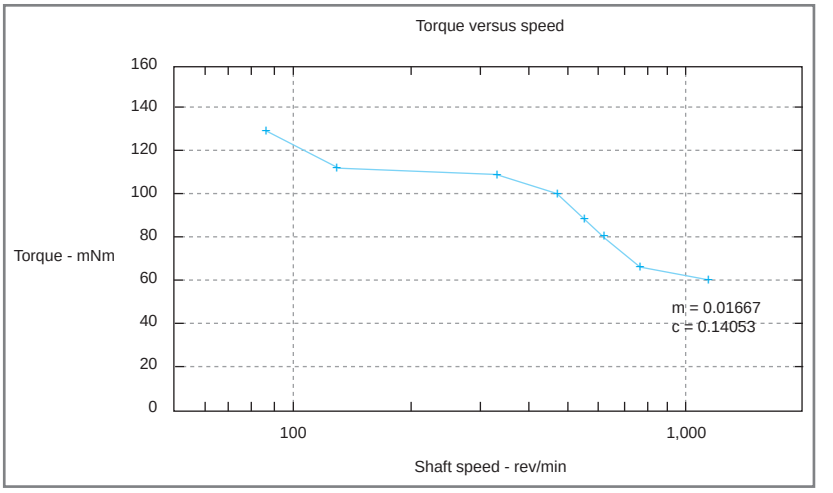


Figure 2—These data points show Minebea's 17PM-J034 motor's maximum torque at various speeds in 1/8 microstep mode with 1 A peak current. The logarithmic X axis compresses the large speed range: equal speed ratios have the same size.

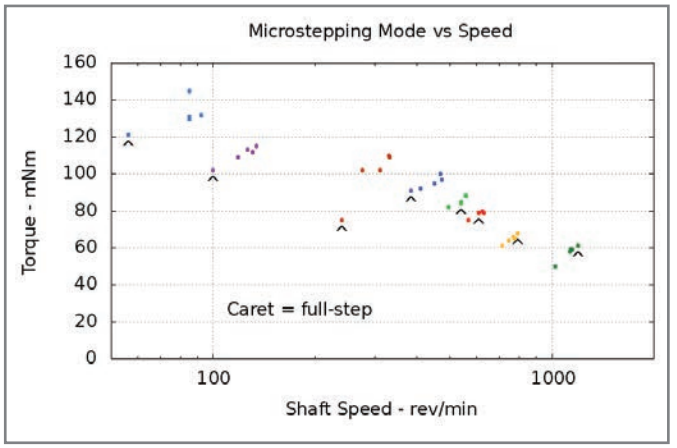


Figure 3—Each Allegro A4988 stepping mode produces a slightly different torque-speed combination for a given load resistance. Full-step mode does not consistently produce more torque than the microstepping modes.

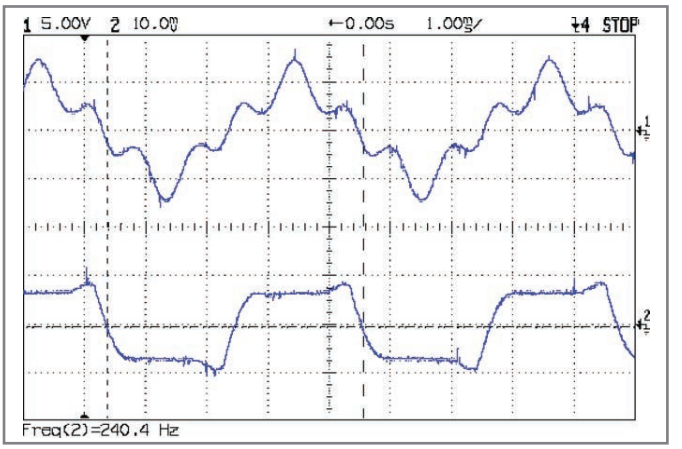


Photo 3—Running at 960 step/s in full-step mode excites a 1,200 Hz mechanical resonance. The upper trace shows the generator's open-circuit output voltage and the lower trace is the motor drive current at 1 A/div.

winding current to the motor. The left-most point came from measuring the current through a dead short applied to the generator terminals, while the others measured the voltage across power resistors as described earlier. The right-most point came from a pair of 50-Ω resistors.

The motor produces reasonably constant torque up to about 500 rpm, then begins its expected torque-versus-speed drop. The last data point seems somewhat too high to continue the downward slope, but that's the real measurement!

TORQUE VS. WINDING CURRENT

I tested the motor with 1-A and 500-mA winding currents to test the often-cited "torque is proportional to current" mantra: half the current should produce half the torque. In 1/8 microstepping mode with 1-Ω load resistors, the torques were 112 mN·m at 131 rpm and 45 mN·m at 36 rpm, with both speeds in the motor's "constant torque" range.

The torque ratio is 2.5:1 or 1:0.40, rather than the expected 2:1 or 1:0.5. However, the motor speed differs by nearly a factor of four, which affects both the microstepping waveform and torque. In retrospect, I should vary the resistor inversely as the square of the current to maintain a more-or-less constant stall rpm.

It's reasonable to use proportionality as a rule of thumb, but without expecting an exact relationship. As always, you should measure your devices at their actual operating conditions.

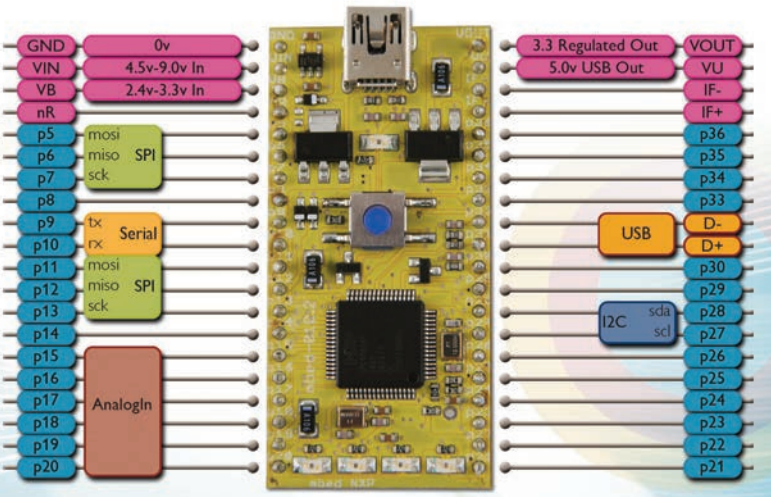
TORQUE VS. MICROSTEPPING MODE

Because motor torque varies linearly with winding current, the reduced current imposed by a microstepping driver should decrease the available motor torque. Figure 3 shows the results of measuring the stalling torque in each of the A4988 step modes (full-step and 1/2, 1/4, 1/8, and 1/16 microstep) with the colors corresponding to eight different load resistances ranging from zero (a dead short) to 50 Ω.

A caret marks the full-step mode within each group of five points. Only

mbed

mbed NXP LPC11U24 Microcontroller



Rapid prototyping for USB Devices, Battery Powered designs and 32-bit ARM® Cortex™-M0 applications

http://mbed.org

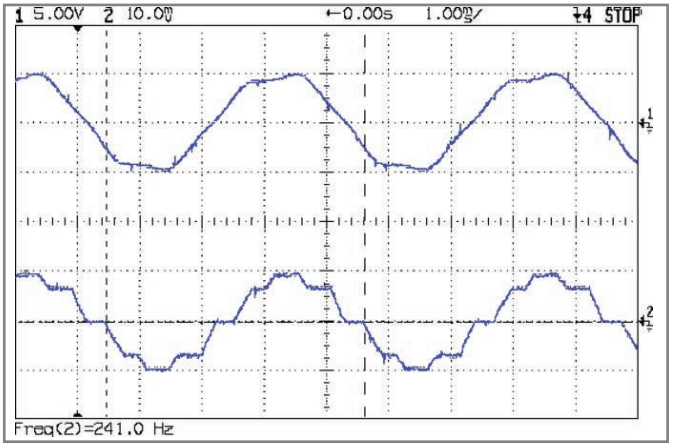


Photo 4—Running at 1,920 step/s in half microstepping mode eliminates much of the 1,200 Hz mechanical resonance, although the generator output voltage shows considerable distortion.

for speeds around 1,000 rpm does full-step mode produce higher torque than the microstepping modes and, in any event, it's well inside the error range.

You can examine the raw data by downloading the spreadsheet for this column.

RESONANCE VS. MICROSTEPPING

The rigid metal-and-plastic dynamometer structure provides a stiff mechanical configuration that's ideal for demonstrating mechanical resonances. I set the DIP switches for full-step mode and adjusted the step rate to the middle of the ear-shattering resonance shown in Photo 3.

The winding current in the lower trace has a 240 Hz frequency (i.e., 960 full step/s = 4.8 rev/s at 200 step/rev). The motor current, a square wave with a peak of 0.7 A, is what the A4988 produces in full-step mode: both windings energized simultaneously at $1/\sqrt{2}$ of the 1 A peak setting. The small bumps just before the transitions show that the A4988 isn't completely controlling the current, but it's doing a reasonable job in the face of the powerful resonance.

The generator open-circuit output voltage in the top trace shows the rotor oscillates at five cycles during the four full steps in each current waveform cycle, imposing a 1,200 Hz resonance on the 240 Hz fundamental waveform. The vibration loosened the setscrews in the steel sleeve connecting the two motor shafts in about a minute; you must use thread locking adhesive in a permanent installation!

Switching to half microstep mode doubles the step rate to 1,920 step/s and moves most of the energy away from the resonant frequency. The generator output in Photo 4 shows some distortion, but the motor runs with far less audible noise and vibration.

Running in quarter microstep mode produces the much cleaner result in Photo 5. One consequence of microstepping is the increased step frequency for a given rotational speed: producing accurately timed steps poses a challenge for microcontroller firmware. If you're doing this sort of thing, figure out how to use on-chip hardware timers for pulse generation!

Photo 6 shows that microstepping doesn't solve all mechanical resonance problems. Running at 960 step/s in half microstepping mode provides enough energy to excite the same howling 1,200 Hz resonance, although the frequency of successive cycles varies from below 1,000 Hz to about 1,300 Hz.

CONTACT RELEASE

You can certainly improve the simple dynamometer shown here, although even in this form, it can provide useful data. You'll find that motor datasheet specs differ from any measurements you can make and both of those differ from what your project actually requires.

Although stepper motors seem to be digital electrical devices, they may drag you deep into mechanical design before your project goes out the door!

Ed Nisley is an EE and author in Poughkeepsie, NY. Contact him at ed.nisley@ieee.org with "Circuit Cellar" in the subject to avoid spam filters.

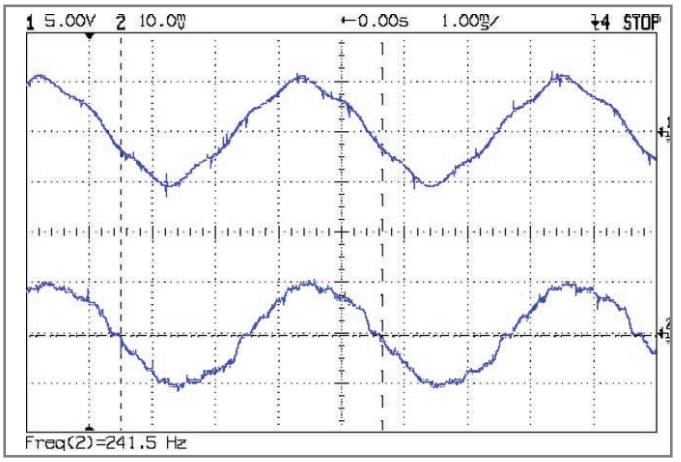


Photo 5—At 3840 step/s in 1/4 microstep mode, the generator output in the top trace shows only a hint of the 1,200 Hz mechanical resonance.

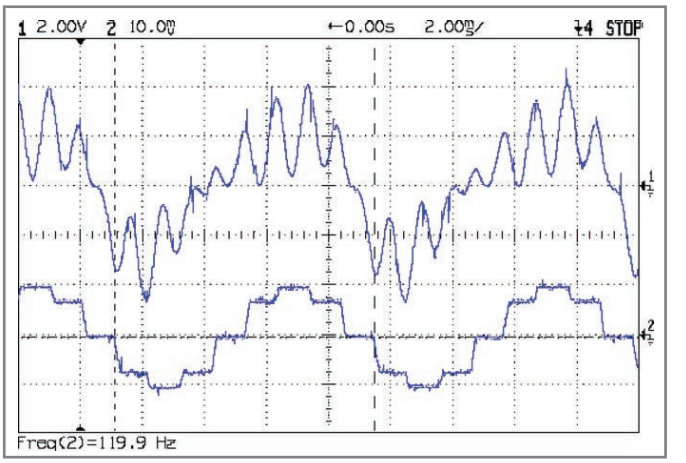


Photo 6—Running at 960 step/s in 1/2 microstepping mode excites the same 1,200 Hz mechanical resonance, but with irregular cycles driven by successive microstep transitions.

PROJECT FILES

To download measurement data and scripts, go to ftp://ftp.circuitcellar.com/pub/Circuit_Cellar/2011/257.

RESOURCES

E. Nisley, "Stepper Motor Drive (Part 1): The Basics," *Circuit Cellar* 253, 2011.

———, "Stepper Motor Drive (Part 2): Microstepping," *Circuit Cellar* 255, 2011.

SOURCES

Allegro A4988 DMOS Microstepping Driver
Allegro MicroSystems, Inc. | www.allegromicro.com/en/Products/Part_Numbers/4988

ACT 17HS5425 Stepper motor
Changzhou ACT Motor Co., Ltd. | www.act-motor.com/product/nav_4_6.html

17PM-J034 and 17PM-J049 Hybrid stepping motors
Minebea Co., Ltd. | www.eminebea.com/content/html/en/hybrid_list/index.shtml

A4988 Stepper motor driver carrier
Pololu Corp. | www.pololu.com/catalog/product/1182

Current probe
Tektronix, Inc. | www.tek.com/products/accessories/current.html

NEED-TO-KNOW INFO

Knowledge is power. In the computer applications industry, informed engineers and programmers don't just survive, they *thrive* and *excel*. For more need-to-know information about some of the topics covered in this article, the *Circuit Cellar* editorial staff recommends the following content:

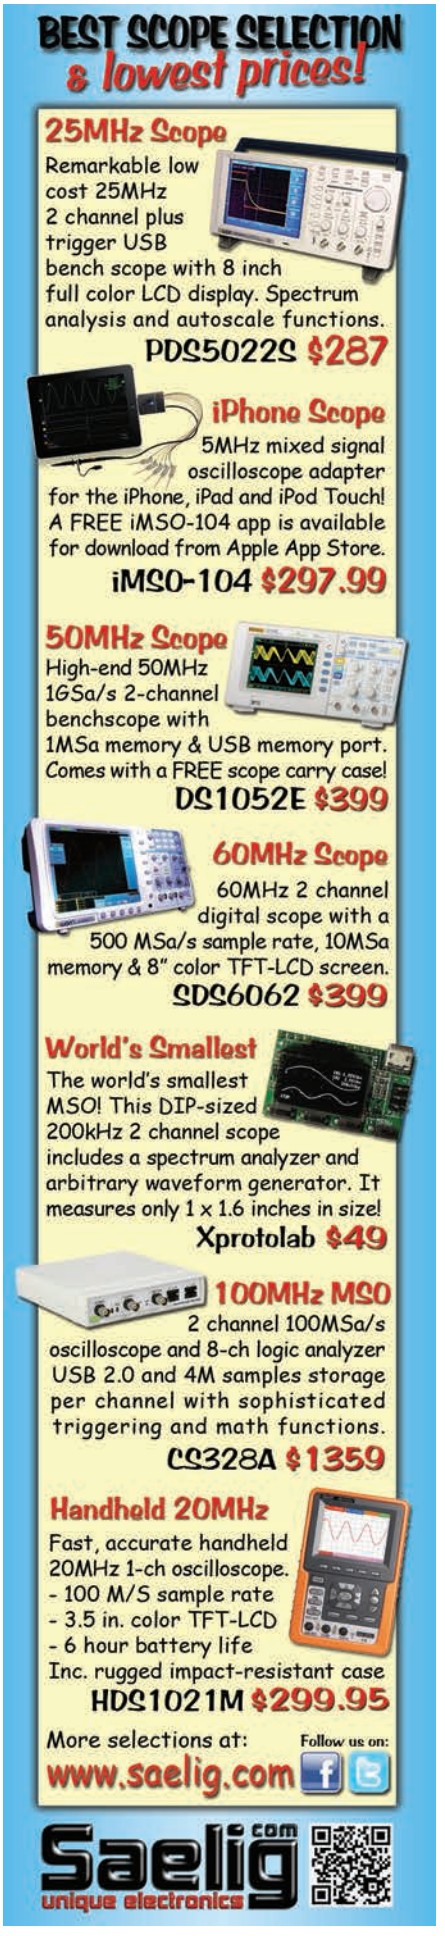
Pump Control
Build a Variable-Speed Drive for Motors
by Richard Wotiz
***Circuit Cellar* 215, 2008**

Richard's single-phase, variable-speed drive for AC induction motors is an excellent device for powerful, yet quiet, pump operation. Designed for use with a capacitor-start/capacitor-run motor, it includes active power factor correction (PFC) and inrush current limiting. Topics: Pump, Motor, Speed Control, AC Induction, PFC, Power Factor, Keypad, Temperature, PIC16F687, MCP4011, LCD

Muscle for High-Torque Robotics
by Lonne Mays
***Circuit Cellar 153*, 2003**

Is your favorite robot getting to be a bit too atrophied for your taste? Does it lack vim and vigor? If so, don't worry. What you need is high torque and speed. With the MC33887D integrated H-Bridge and a little direction from Lonne, you can bulk up your robot in no time. Topics: Robotics, Torque, MC33887D, H-Bridge, Speed, Servo Amp, Analog Switch, Test-Pulse Generator

Go to *Circuit Cellar*'s webshop to find these articles and more: www.cc-webshop.com



by George Novacek (Canada)



EMI and Data Integrity

Electromagnetic interference (EMI) is the most common cause of electronic data corruption. While error-correction algorithms are frequently used to protect data integrity, most increase data redundancy, which also increases bandwidth and cost. As you'll see, shielding can protect against EMI and minimize the possibility of data corruption by increasing system immunity to interference.

Data integrity is essential for making decisions. Its absence can be characterized by the old computer adage "garbage in, garbage out." To protect data integrity, many interface protocols contain error-correction algorithms. Unfortunately, error-correction schemes require an increase in data redundancy. That means more bandwidth, and there's a cost associated with it. It is, therefore, preferable to minimize the possibility of data corruption by increasing system immunity to interference, as it makes the work of the error-correction mechanism more reliable and data redundancy can be decreased, for example, by stronger compression.

ELECTROMAGNETIC INTERFERENCE

The most prevalent cause of electronic data corruption is electromagnetic interference (EMI). It is distributed through the air by radiation, through interconnecting wires by conduction, and most often by a combination of both. The sources of interference can be climatic conditions, or electrical equipment such as radio and TV stations, radars, computers, electric motors, or engine ignitions.

Shielding is the first step in protecting data from EMI. In my last column, I discussed transfer impedance as it is related to the effectiveness of shielding of wires and cables ("Shielding and Transfer Impedance," *Circuit Cellar* 256, 2011). At the time, for simplicity, we considered the EMI to be injected into the shield by a current transformer. But this is hardly a real-life situation. I can't imagine that, other than in a test lab, interface harnesses were purposely run through such a transformer. But it is a convenient method to simulate the effects of the interference and to test susceptibility of the equipment to frequencies from 10 kHz up to about 400 MHz. One reason for this is that efficient antennas for frequencies below 100 MHz are too large to be practical for installation in anechoic test chambers. Only at 100 MHz antenna size becomes practical to work with.

Another reason is that the electromagnetic field behavior falls into two distinct areas. From about two wavelengths (2 λ) distance from the antenna to infinity exists what is called the far field. There are several theories as to precisely where the far field begins. But, if the tested equipment is more than twice the wavelength away from the radiator, it is in the far field. Up to one lambda (λ) is near field, and a 1- to 2-λ distance is considered the transition field.

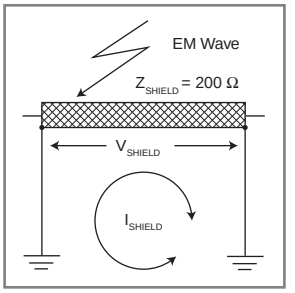


Figure 1—Electromagnetic wave induces current in the shield

As the name implies, an electromagnetic field comprises two components: electrical (e-field) and magnetic (h-field). The h-field is prevalent in the near field; the e-field dominates the far field. Most EMI happens in the far field, where the wave behaves in a "normal," predictable fashion, such as the victim having no effect on the



▶ CUWIN6300
▶ CUWIN5500
▶ CUWIN5200

Windows CE 6.0 Touch Controller

CUWIN

The CUWIN is a series of Windows CE touch controllers that are more cost-effective than a PC, but with more features than an HMI touch screen.
Create sophisticated applications with C++ or any .Net language.

533MHz ARM CPU
128MB SDRAM & Flash
SD Card Support
Ethernet, RS-232/485
USB, Audio Out
Windows Embedded CE 6.0



C-Programmable Modular Industrial Controller

MOACON

The MOACON is a modular, C programmable, ARM-based automation controller designed for industrial environments.

Choose from a diverse, feature-rich selection of modules including:
Digital I/O
Analog I/O
RS-232/485
Motor Control
Ethernet
High-speed Counter & PWM

BASIC with LADDER LOGIC CONTROLLER



New

Integrated CUBLOC Controller and I/O Board

CB210

The CB210 is an inexpensive, integrated CUBLOC controller and I/O board programmable in both BASIC and Ladder Logic.

I/O Ports x20	10-bit A/D x6
PWM x3	RS-232 x1
80KB Program Memory	3K Data Memory

The CUPC is a series of industrial touch panels with all the features of a modern PC for the most feature-rich user experience.

ATOM N270 1.6GHz CPU
2GB RAM
320GB HDD
SD Card Support
Color Display (1024 x 768)
RS-232, Ethernet
USB, Audio Out
Windows XPe/XP Pro

CUPC

Industrial Touch Panel PC with ATOM Processor



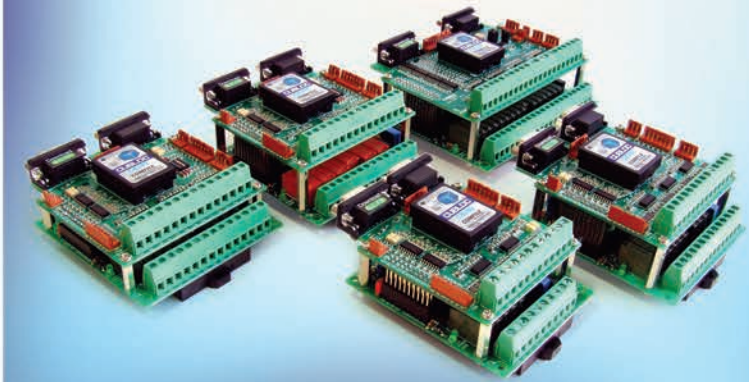
▶ CUPC-P200
▶ CUPC-P220

The CUSB is a series of compact, CUBLOC-integrated, industrial I/O boards programmable in both BASIC and Ladder Logic.

CUBLOC CB280 Core Module
80KB Program Memory
3KB Data Memory
DC 24V Inputs x9~16
A/D Inputs x2~6
Relay Outputs x6~16
PWM x2~6
High-speed Counters x0~2
RS-232/485 x1~2

CUSB

Integrated Industrial Controllers



COMFILE TECHNOLOGY
1175 Chess Dr., Suite F, FOSTER CITY, CA 94404
call : 1-888-9CUBLOC (toll free)
1-888-928-2562
email : sales@comfiletech.com
www.comfiletech.com

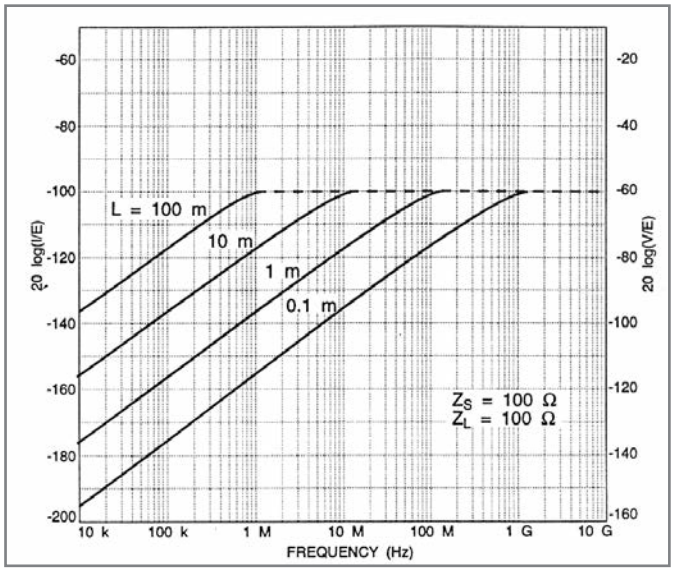


Figure 2—Voltage induced in a length of wire in 1-V/m e-field

source and the power diminishing with the square of the distance from the source. I shall address h-field issues in a future column.

Consider testing the effects of 1-MHz EMI. Its far field starts 600 m (1,835′) from the antenna, making the test setup impractical. The transformer injection solves the problem. Above 100 MHz, the 6-m (20′) distance from the antenna becomes reasonable, so the wiring harness can be illuminated by an e-field, whose strength is expressed in volts per meter. The system designers are responsible for the definition of the field strength to be encountered in the operating environment.

SAFETY & SHIELDING

Most consumer equipment is immune to 1 V/m to 2 V/m. Such immunity can be achieved without much special effort, other than good layout techniques and common sense. Safety-critical requirements—such as those for medical, automotive, and aerospace equipment—have, over the years, gradually increased. Where 30 V/m immunity was not long ago considered adequate, 200 V/m is normal by present-day standards. Some frequency bands, usually above 1 GHz, require immunity measured in thousands of volts per meter. With interfering frequencies from 100 MHz to 20 GHz, sometimes up to 40 GHz, just 200-V/m EMI suppression is no easy task. At 40 GHz, the wavelength is mere 7.5 mm (0.29″). A quarter-λ discontinuity (that is 1.88 mm or 0.07″) in the shielding or the enclosure can act as a resonator or a slot antenna, making the equipment susceptible at that band.

How do we proceed designing the equipment when the customer specifies the EMI field strength to which the equipment must be immune? Let's look at Figure 1, which shows the basic assumption of the shielding exposure to EMI. Figure 1 is a modified diagram from my article "Shielding and Transfer Impedance." We are now concerned with the voltage developed by the EM wave illuminating the shield. Let's further assume the shielded interface cable is 10 m (32′) long and the EMI field strength 200 V/m.

Figure 2 shows an asymptotic relationship between the frequency and the cable length normalized for illumination by 1 V/m field. The shield of the shielded cable is assumed to be located 1.4 mm (0.05″) above ground, which is essentially the insulation thickness. In Figure 2, we can also see that from 10 MHz up the induced voltage in the 10-m cable shield will be constant. Below 10 MHz it will decrease at a rate of 20 dB per decade.

On the right axis of the diagram, we determine that the induced signal above 10 MHz is –60 dB referenced to 1 V/m field. Typically, system harnesses are tested at 50 mm (2″) above ground plane, which is a safe assumption for an actual installation. Looking at the correction diagram shown in Figure 3, we interpolate the necessary correction to be 26 dB. That means that the voltage induced into the shield will be –34 dB (i.e., –60 dB + 26 dB). At 1 V/m field strength, the voltage, corresponding to V_{SHIELD} in Figure 1, will be 20 mV. Shield impedance of 200 Ω is assumed, that is $Z = Z_S + Z_L$ per Figure 2, distributed along the shield. The 200-V/m e-field will thus induce 4 V_{SHIELD} (i.e., 20 mV × 200). The left axis in Figure 2 indicates the shield current at that impedance. For different impedance and field strength other than 1 V/m, we can calculate the shield current using Ohm's law.

Figure 3 also illustrates how important it is to run cables close to the ground (or the chassis). If the cable is 1,400 mm (4.6′) above the ground, the induced baseline voltage derived from Figure 2 will increase more than 1,000 times (60 dB). Thus, a 50 mm (2″) maximum harnesses height above the ground plane should be the system installation worst-case condition.

CALCULATING INTERFERENCE

Having established the shield current, the transfer impedance of the cable, and the load impedances of the shielded conductor (refer to my November article), we are

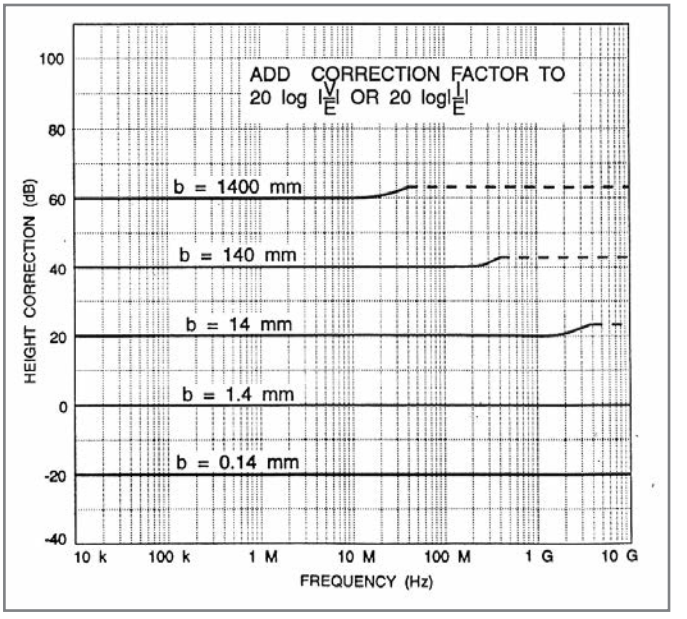


Figure 3—Correction factor based on installation

able to calculate the interference signal magnitude to be seen by the equipment.

The very same calculations of interference levels for immunity from EMI apply for calculating the EMI generated internally by the equipment and emitted to the outside world. The source and the victim are merely reversed.

With the magnitude of the EMI seen by the equipment and vice versa known, we can design the topology and shielding of the interface circuits to effectively control it. But we must always remember that the calculations are not precise due to many environmental factors and a healthy safety margin must be included. I shall address the protection circuits and the effects of h-field in future columns.

George Novacek (gnovacek@nexicom.net) is a professional engineer with a degree in Cybernetics and Closed-Loop Control. Now retired, he was most recently president of a multinational manufacturer for embedded control systems for aerospace applications. George wrote 26 feature articles for Circuit Cellar *between 1999 and 2004.*

RESOURCES

G. Novacek, "Shielding and Transfer Impedance," *Circuit Cellar,* 256, 2011.

O. Hartal, *Electromagnetic Compatibility by Design*, R&B Enterprises, 1993.

NEED-TO-KNOW INFO

Knowledge is power. In the computer applications industry, informed engineers and programmers don't just survive, they *thrive* and *excel*. For more need-to-know information about some of the topics covered in this article, the *Circuit Cellar* editorial staff recommends the following content:

Error Checking
by Jeff Bachiochi
***Circuit Cellar* 250, 2011**

Data security isn't a perk. It's a requirement. This article covers the topics of error checking, checksums, and the cyclic redundancy check. Error checking matters. Topics: Error checking, Checksum, Cyclic Redundancy Check, CRC

Cable Shielding Experiments
by Robert Lacoste
***Circuit Cellar* 219, 2008**

After reading about the topic of noise coupling with various cables and wiring setups, Robert reproduced the cable shielding experiments described in the book. This month, he shares the results. Topics: Shielding, Noise Reduction, Filter, Capacitive Coupling, Grounding, Signal Source

Go to *Circuit Cellar*'s webshop to find these articles and more: www.cc-webshop.com

Get PUBLISHED. Get NOTICED. Get PAID.

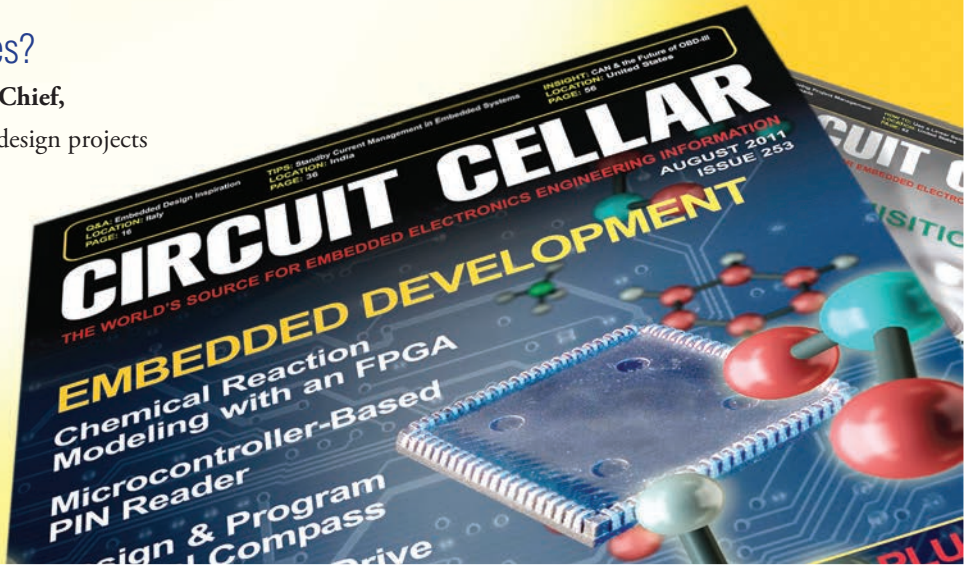
Circuit Cellar feature articles are contributed by professional engineers, academics, and students from around the globe. Each month, the editorial staff reviews dozens of article proposals and submissions. Only the best make it into the pages of this internationally respected magazine.

Do you have what it takes?

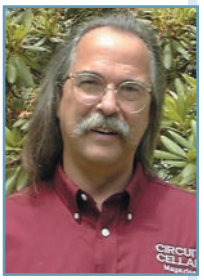
Contact **C. J. Abate, Editor-in-Chief,** today to discuss the embedded design projects and programming applications you've been working on and your article could be featured in an upcoming issue of *Circuit Cellar* magazine.

editor@circuitcellar.com

CIRCUIT CELLAR



by Jeff Bachiochi (USA)



Fly-By-Wire Wheelchair (Part 2)

Bluetooth Wheelchair Control

The first part of this article series covered the analog joystick design in a power wheelchair. This article details how to incorporate wireless connectivity and a PC application to display joystick positioning.

Many great electronic designs are solutions to problems engineers have experienced firsthand. This article details how I came work with a power wheelchair and determine that I could use my engineering skills to build an advanced control system.

For the most part, a power wheelchair can give those with temporary or permanent mobility issues a way of interacting with the world in a more personal way. But, for people with more restricted motion, steering a power wheelchair via a joystick may not be possible. As one might suspect, many wheelchair manufacturers aren't prepared to take on the special needs of this smaller group of individuals. But an engineer can remedy this situation.

You can find a number of input devices to support the special requirements of people who can't easily operate a joystick by hand. But, like the luxury car add-ons, they are usually high-cost and available on only a few models. This project makes these kinds of devices accessible to a standard power wheelchair that doesn't offer such a level of utility.

In the first part of this series, I described the essence of the project—the analog joystick used by most manufacturers. Now I'll explain how I provided the control system with my interpretation of the analog joystick using an analog-to-digital converter (ADC) to measure the present joystick position and a digital-to-analog converter (DAC) to reproduce it. As you'll learn, once that was accomplished, I was able to hijack the DAC and provide control from an external source.

The circuit I presented in the first article provides a serial universal asynchronous receiver/transmitter (UART) connection for some external circuitry to take control. To make this project more flexible, I added a Bluetooth transceiver for wireless connectivity (see Photo 1). I finished up by defining a simple control protocol and actually sending converted joystick data out the serial port. A PC application was written in Liberty BASIC display joystick position as a graphic.

EXPERIENCE LEADS TO INSIGHT

Once a year, my boy scout troop has a Disabilities Night. When the boys enter the meeting hall, they are assigned a random disability by picking from a hat. For the rest of the meeting, they must function as though they were handicapped. To help them feel the part, they are given some assistance—or should I say hindrance. Some receive a blindfold, crutches, or earplugs. Others receive goggles smeared with petroleum jelly or some other sense-altering paraphernalia. Typically they pair up with others who have a different

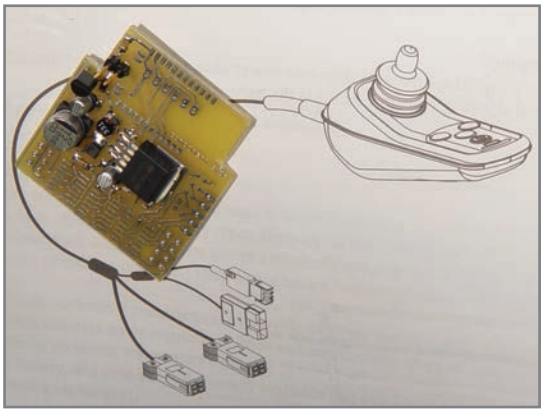


Photo 1—The SMT prototype shows the real estate relationship between the populated power supply section and the rest of the unpopulated circuitry.

disability and work together to provide assistance to each other.

The meeting's agenda is similar to most, so the boys are already familiar with many of the activities. It's a good combination of learning and fun. However, on this night, they take home much more than a new scouting skill. We spend a few minutes before the closing to let each boy talk about his experience. This sharing helps all to understand the difficulties of not just their own temporary handicap, but also all other disabilities.

I find there is nothing like experience to help one understand things more clearly. With access to a power wheelchair for the project I'm describing in this article series, I've had the opportunity to experience its benefits firsthand. While the front door to my house is a mere one step up, the back door (and the one most often used) is five steps. In case you haven't realized, wheelchairs were not designed to go up and down stairs. So, I was not able to get into my house while in the wheelchair. Both my bedroom and office are on the second floor. Without a lift, I wouldn't be visiting either of these. A couple of 2 × 4s and a sheet of plywood enabled me to get up and into our kitchen (which is located off the front door). I don't know about your house, but mine wouldn't make a magazine spread, unless it was for "*Clutter Illustrated.*" It's a constant battle between putting things away and putting them down. Needless to say, there are not straight open pathways between rooms.

TOUR DE FORCE

The joystick controls made it easy to move the wheelchair. From a technical point of view, I found forward and reverse to be as expected, while left and right were more complicated. Pushing the joystick to the left (or right) rotated the wheelchair on its axis to the left (or right). While this seems natural, think about what is happening with the wheels. They must move in opposite directions. While this is also predictable, it is the motion of the wheels from all forward (or reverse) to all left (or right) that is a bit unexpected.

With a 90° joystick sweep from forward to left, the left wheel goes from forward to stop to reverse, while the right wheel remains constantly driving forward (assuming a constant speed). With the joystick at approximately 45°, the left wheel isn't moving at all! This extra complexity is required to enable the wheelchair to rotate upon itself for the tightest turning radius possible. The drive wheels measure 23" (wheel to wheel), which enables the wheelchair to fit easily between doorways. However, the actual turning radius is about twice that because of the outrigger wheels located in the back corners of the frame. Because the frame is essentially rectangular, as opposed to round, the minimum turning radius is a diagonal and not just a diameter.

What this means for accessibility is that while doorways themselves are not obstacles, the areas around them are problematic, especially where a turn is involved. I found I had to do a multipoint turn (like trying to turn a car around in a narrow roadway) in order to maneuver between rooms. Luckily, we have a bathroom on the first floor. Getting into the bathroom was easy enough, but using any of the fixtures was next to impossible.

I discovered the aforementioned problems after using the chair for only a short period of time. I'm sure I'd find many more issues if I were to test it in various other locations and environments. But what I learned was enough to get the gears in my mind turning. Let's now focus on the engineering.

> "The idea here is to provide the control system with my interpretation of the analog joystick using an ADC to measure the present joystick position and a DAC to reproduce it. Once this was accomplished, I could then hijack the DAC and provide control from some external source."

RX DATA

Five bytes of data are sent in an 11-byte packet consisting of five hexadecimal values (ASCII pairs) followed by a <CR>. These values are the most significant 8-bits of each of the A/D conversion results for the four joystick channels and a control byte. At this point, the control byte reflects the status of the "emergency input" (the only digital input to my circuit) as a 1 (not pushed) or a 0 (pushed). The PC application has a Loopback mode that will send back data as received. This isn't practical in a logical sense, as Rule 2 of the micro code is: "All external data is discarded if the onboard joystick is not idle (centered)." So, as soon as the joystick is moved in any direction, data to the DAC comes from the onboard joystick. However, all external data is received, and if the packet is of legal size, the data is stored for use if the onboard joystick is idle. I can check this data against what was sent using the in-circuit debugger (ICD), just to make sure the RX routine is working properly.

The "emergency input" satisfies Rule 1: "Stop everything if emergency input is grounded (pushed)." This causes idle data to be fed to each DAC, which will disable motor current and apply the parking brakes.

Assuming neither rule is in effect—that is, the emergency input is open and the onboard joystick is idle—we can enable external data to be transferred to the DACs. This data is coming from the PC application (in this case) but could come from any external device.

I took a USB flight simulator joystick and added it to my PC application. I can redirect its X and Y values to the application's joystick display in place of the onboard joystick data that's being sent via Bluetooth. I want to make sure I have transposed the flight simulator's joystick axis correctly and that the resulting data is the same from each input device.

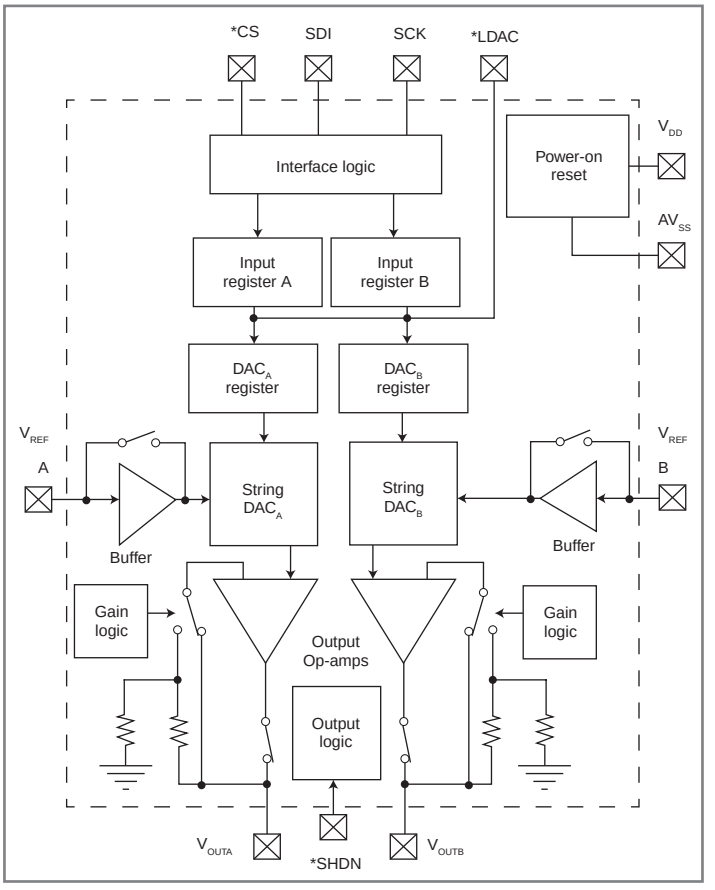


Figure 1—Microchip Technology's MCP4922 dual-channel DAC comes in three varieties: 8-, 10-, or 12-bit resolution. Each has identical packaging and command formats, which makes upgrading or downgrading a simple physical exchange.

With the flight simulator's joystick output equal to that of the onboard joystick, I feel comfortable substituting the flight simulator's data for that of the onboard joystick. Pull-down menus in the PC application enable the selection of both an input device to the joystick display on the PC and a source for the data sent back via Bluetooth to the project circuitry. I think I'm ready to go for a live test. But first…

(R2)D2A

I chose a pair of Microchip Technology MCP4922 dual-channel DACs to mimic the four voltage outs that normally come directly from the onboard joystick. Figure 1 is a diagram of the converters. These DACs come in 8-, 10-, and 12-bit resolutions. While I've decided to base this project on 8-bits, its nice to be able to get a higher resolution just by swapping out chips. These have a buffered output and can source greater than 10 mA. If you recall from last month's discussion, the onboard joystick has dual-voltage outputs (X1, X2, Y1, and Y2). I could have used a single D/A chip and fed the X and Y output to both X1 and X2 and Y1 and Y2 inputs to the VR1 (wheelchair controller), but I wanted to retain its safety-related function. That is, if the VR1 doesn't see equality between X1 and X2 (or Y1 and Y2), then it doesn't trust the outputs and puts the wheelchair in a safe mode. Note: This rule will also be applied to any external data and, should it receive corrupted data (X1 doesn't match X2 or Y1 doesn't match Y2), the VR1 will consider this an unsafe condition and "safe" the wheelchair.

Data is sent to the DACs via a SPI. Separate CS lines enable each chip to be selected independently. Sixteen bits of data are transferred for each channel. That's 4 bits of configuration and 12 bits of data. The data is left-justified so the transmission remains the same (16-bits) for the 8-, 10-, and 12-bit parts! There are separate V_{REF} inputs for each channel (see Table 1).

HOT SEAT

With circuit in hand, I sat down in the power wheelchair and connected it up between the on-board joystick and the VR1 in the armrest. A push of the Power button, and the LEDs began dancing around indicating an error message. According to the manual, the joystick wasn't centered on power-up. I went back into my code and forced an idle condition upon initialization, but I got no further with that. Out came the voltmeter, and I found the 5-V power to be low. My circuit seemed to be putting too much of a load on the VR1 circuitry (a maximum of less than 100 mA). Talk about designing close to the edge! Maybe measuring for an overcurrent was a safety issue programmed into the VR1's operating system? I decided to try supplying power to my circuit externally with a wall-wart power supply. This time, the power wheelchair started up with no errors. I gingerly tried the on-board joystick and I had normal wheel movement. Great. So far, so good. But I wasn't going to move very far tethered to the outlet holding the wall-wart supply.

The VR1 controller is powered from the battery connections (two 12-V batteries in series). I could get the circuit power I needed from the raw 24-V inputs. One hundred milliamps dropped across a 7805 regulator would be 1.9 W (i.e., 24 V – 5 V = 19 V × 100 mA). Yikes, that's a lot of heat to dissipate here! So, I looked for a switching regulator that would eliminate this wasted energy and, more importantly, unwanted heat.

A Texas Instruments TL2575-5 voltage regulator was a good fit. The fixed-output voltage is regulated by keeping

Bits															
Control				Data											
15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
Channel select	Output buffer enable	V_{REF} 2× enable	Output Enable	D11	D10	D9	D8	D7	D6	D5	D4	D3	D2	D1	D0

Table 1—This is the SPI communication format for talking to a Microchip Technology MPC49x2 dual-channel DAC.

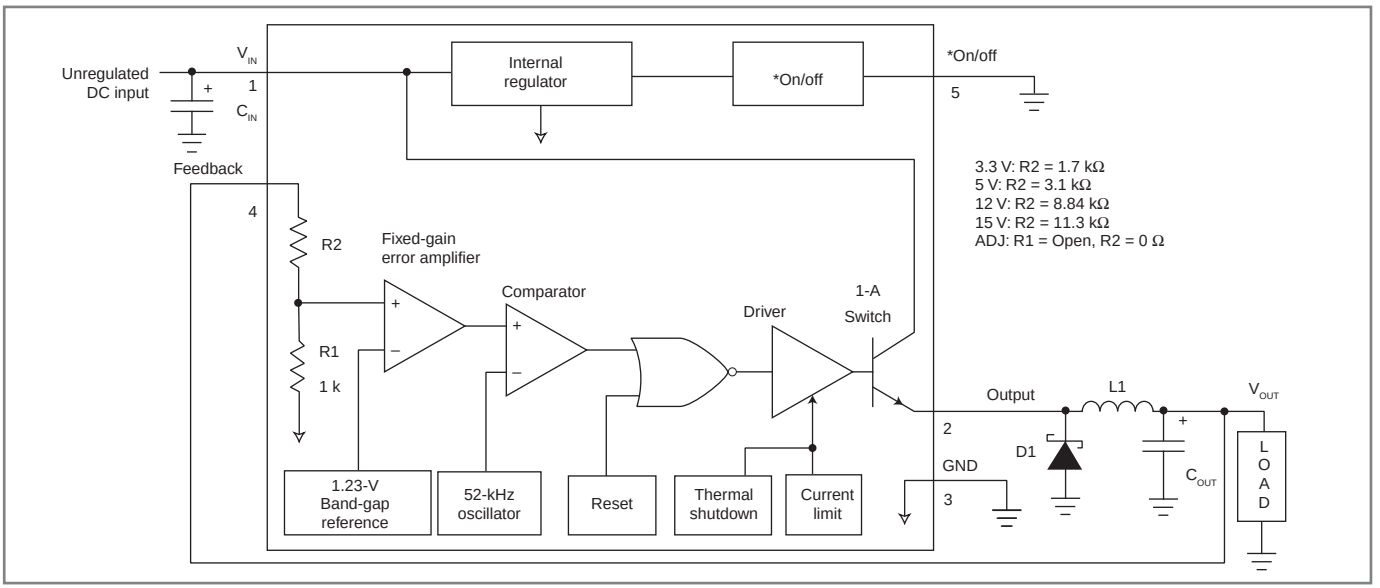


Figure 2—The Texas Instruments TL2575-5 switching regulator can reduce the normal loss across a linear regulator by applying an input only when necessary to sustain a fixed-output voltage.

the output capacitor (C_{OUT}) charged properly via the inductor (L1) connected to the supply via the internal pulse-width modulation (PWM) circuitry (see Figure 2). The cycle-by-cycle control adds some digital noise, but this is held to approximately 50 mV with no additional filtering.

To prevent any voltage outs from the DACs from affecting any unpowered circuitry, I connected the shutdown inputs on the MCP4922s to the VR1's 5-V output on the eight-pin connector (SV2). This will disable any output unless there is power to the on-board joystick that I've hijacked. Unfortunately, the TL2575-5 uses a logic-low ON signal to enable its operation, or I could have used the VR1's 5-V output to enable my switching regulator. Anyone know of a switcher with logic-high ON control?



On my second prototype, I was forced into using SMT parts because of the physical size of the switching power supply. Photo 1 shows just the power supply parts built onto the new prototype. A good deal of the total real estate is taken up by the supply and the Bluetooth module (area to the left of the supply).

PC APPLICATION

Back to the wheelchair. I powered on and it looked all right. I tried the on-board joystick to move the wheelchair and everything seemed normal. Next, I powered up the laptop running my joystick application. With the application configured to show the joystick position from the on-board joystick, I positioned the laptop so I could see the screen from the wheelchair and retest the wheelchair movement. The laptop's application mirrored proper joystick movement.

Recently, I added code to this application to gather joystick data from any joysticks plugged into the PC. The last time I used a joystick on my PC, it was analog and it came into the PC through the DA15 game connector. USB has all but taken over for any kind of access to a PC. My USB joystick has a gazillion buttons on it, but I was interested only in the X and Y outputs in this application. The format was a 16-bit number for each axis, so I needed to lop off the low bits and use just the upper 8 bits to make it compatible with the data coming from the power wheelchair circuit over the Bluetooth interface (virtual serial port). By substituting the USB joystick values for that of the received data from hijacking circuitry, I could verify that the motions match those of the on-board joystick.

With one click on a pull-down menu, I sent the USB joystick data out the Bluetooth connection and moved the wheelchair from afar. When I tilted the USB joystick forward, the wheelchair began to move. I then added a little speed and I was off racing toward the TV screen across the room. When I let off on the joystick, the wheelchair came to a halt. I carefully turned the wheelchair around and headed back across the short expanse pushing the joystick to its limits. And then it happened. All the lights on the LED display began flashing as the wheelchair came to a halt.

GOOD PROGRAMMING

I sat down, scratched my head, and pondered this for a while. It seemed like the VR1 did not like something I was doing. Then it hit me. The on-board joystick had a limited output swing of 1 to 4 V. My USB joystick could go from 0 to 5 V. I looked at the difference in joystick display between the maximum excursions of the received data from the on-board joystick and that of the USB joystick. Sure enough, the USB joystick moved the display further because of its higher output range. I wondered: Should I limit its range? I decided the correct way to handle this was in my embedded code. I filtered any axis values (the first four values) received via the serial port, substituting minimums (0x33) and maximums (0xCC), which correspond to 1-V minimums and 4-V maximums. The VR1 was doing its job. It understood that any voltage above or below these levels meant that the on-board joystick (or in this case, the external joystick) was malfunctioning!

UP & ROLLING

Everything was looking good now. I could drive the wheelchair around from afar and, if the on-board joystick was moved, its output took precedence over any external control. The external emergency switch input forced idle values (0x80) to all DAC channels producing idle voltages of 2.5 V on all axis inputs to the VR1. This prevents the on-board joystick or any external device from controlling the wheelchair. In last month's article, I suggested there was a mystery use for this input, other than bringing the wheelchair to a halt in the event of some emergency. Note: An indication of this input's status is in the fifth byte sent to any external circuitry. To an external circuit, this is an indicator that the axis information being transmitted is *not* going to the VR1. (Its inputs are being artificially held at idle.) In this mode, the external circuitry can interpret joystick movements as commands to initiate other activities. I'm not defining anything here, just trying to spark some imagination on what could be done with this knowledge. Mouse input on a PC? IR TV remote? Robotic arm that reaches into the fridge? Think about it.

Is this project's goal of adding a way to remotely control a power wheelchair of any benefit? Perhaps it offers a caregiver a better way of moving a wheelchair without having to reach over from the back or walk alongside it. If your disability included not being able to handle a joystick, how much would a hat with a Bluetooth-enabled accelerometer be worth? Look ma, no hands!

Jeff Bachiochi (pronounced BAH-key-AH-key) has been writing for Circuit Cellar *since 1988. His background includes product design and manufacturing. You can reach him at jeff.bachiochi@imaginethatnow.com or at www.imaginethatnow.com.*

PROJECT FILES

To download the code, go to ftp://ftp.circuitcellar.com/pub/Circuit_Cellar/2011/257.

RESOURCES

Liberty BASIC, www.libertybasic.com.

SOURCES

KC-21.6 Bluetooth data module
KC Wirefree | www.kcwirefree.com

PIC16F1823 Microcontroller and MCP4922 DAC
Microchip Technology, Inc. | www.microchip.com

Jet 3 Ultra motorized power wheelchair
Pride Mobility Products | www.pridemobility.com

TL2575 Simple step-down switching voltage regulator
Texas Instruments, Inc. | www.ti.com

ICbank
Gateway to Electronics

3 PORT INTERFACE
RS-485 to Ethernet Converter



Only
$170



RS-485 to Ethernet Converter

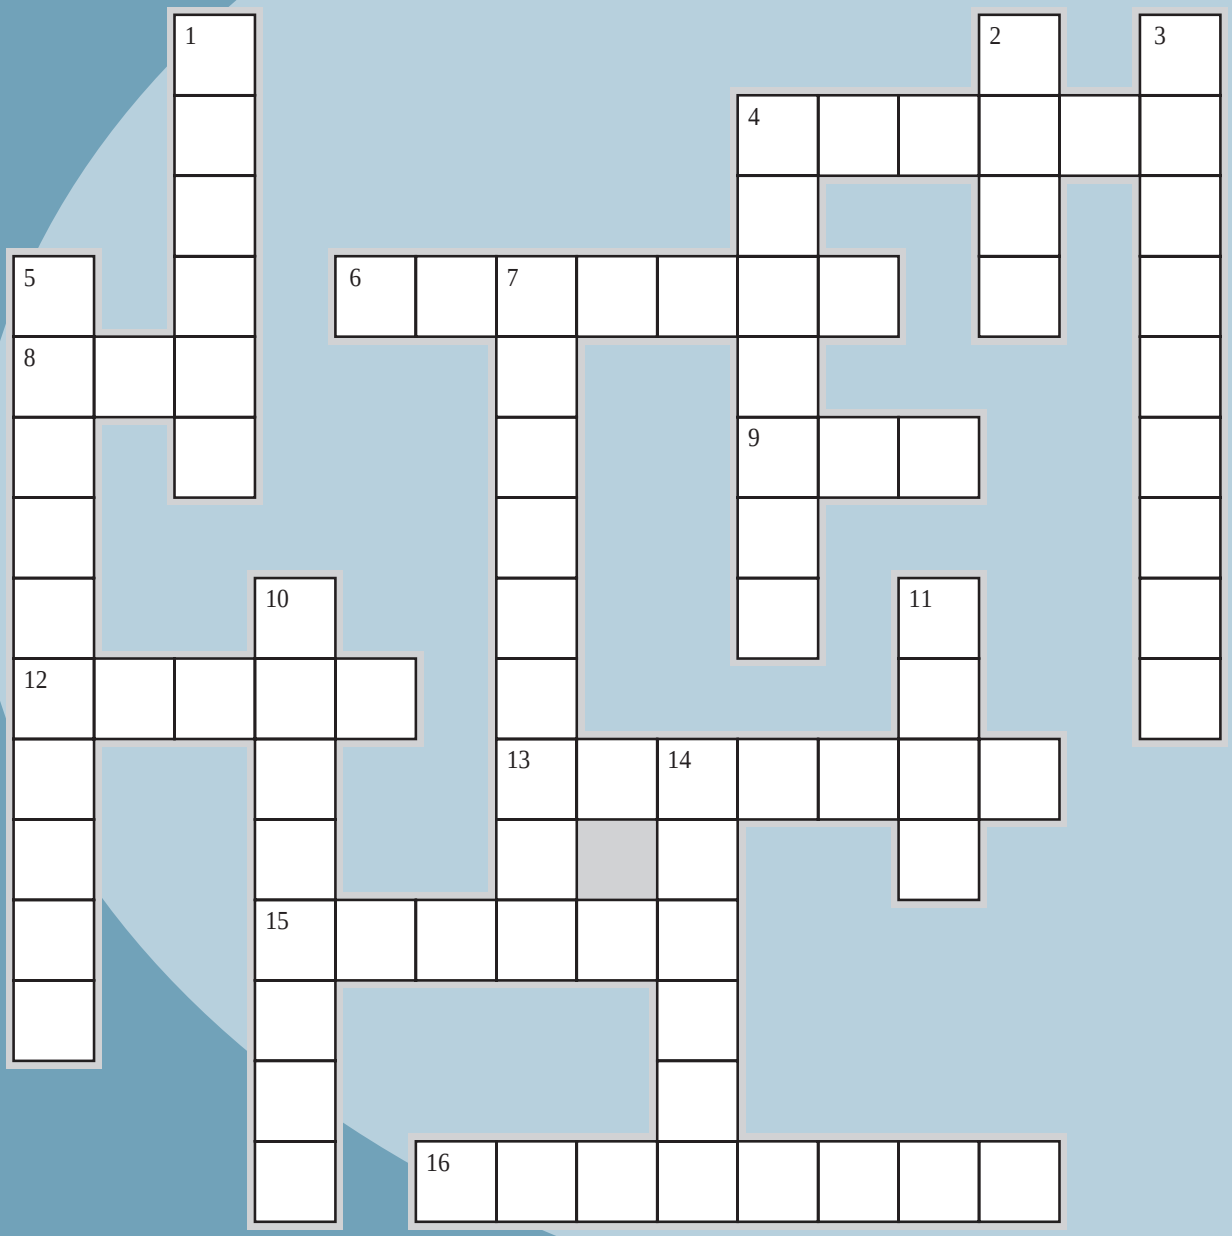
Powerful feature

- Protocol converter RS485 between Ethernet
- Offer TCP/IP Communication to Devices with RS485 I/F

Specification

Network	: TCP, UDP, DHCP, ICMP, IPv4, ARP, IGMP, PPPoE, Ethernet, Auto MDI/MDIX , 10/100 Base-TX Auto negotiation (Full/half Duplex)	
Serial	: RS485 3 Ports, 1,200~115,200 bps, Terminal block I/F Type	
Control program	: IP Address & port setting, serial condition configuration, Data transmit Monitoring	
Accessory	: Power adapter 9V 1500mA, LAN cable	
Etc	: - DIP Switch(485 Baud Rate setting)	- LED: Power, Network, 485 Port transmission signal

CROSSWORD



Across

4. Increase reliability by reducing classification
6. Having not just one polarity
8. A hole in a board
9. Error
12. Offers authentication
13. What happens when you've lost a bit [two words]
15. When combined, more than one forms a product
16. Used for digitial synchronization [two words]

Down

1. AKA, "the perpendicular"
2. Use one data word to alter another
3. A lab analysis simulating an actual I/O [two words]
4. Opposing buses [two words]
5. Caused by too much voltage or current
7. Can make one odd [two words]
10. 000000, for example
11. Right at the heart of things
14. Tau symbol

The answers will be available in the next issue and at www.circuitcellar.com/crossword.

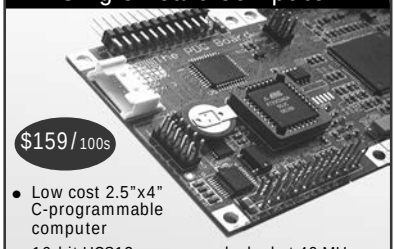
IDEA BOX

THE DIRECTORY OF PRODUCTS AND SERVICES

AD FORMAT: Advertisers must furnish digital files that meet our specifications (www.circuitcellar.com/advertise). ALL TEXT AND OTHER ELEMENTS MUST FIT WITHIN A 2" x 3" FORMAT. E-mail adcopy@circuitcellar.com with your file or send it to IDEA BOX, Circuit Cellar, 4 Park Street, Vernon, CT 06066. For current rates, deadlines, and more information contact Peter Wostrel at 800.454.3741, 978.281.7708 or peter@smmarketing.us.

The Vendor Directory at www.circuitcellar.com/vendor/
is your guide to a variety of engineering products and services.

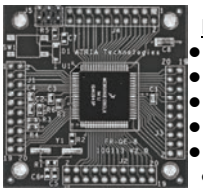
PDQ Board™ - A Fast I/O-Rich Single Board Computer



$159/100s

- Low cost 2.5"x4" C-programmable computer
- 16-bit HCS12 processor clocked at 40 MHz
- 8 PWM, 8 counter/timer, and 8 digital I/O
- 16 10-bit A/D inputs
- Dual RS232/485 ports, SPI and I²C ports
- 512K on-chip Flash, 512K RAM with Flash backup
- Plug-in I/O expansion, including Ethernet, Wi-Fi, GPS, 24-bit data acquisition, UART, USB, Compact Flash card, relays, and more ...

Mosaic Industries Inc.
tel: 510-790-1255 fax: 510-790-0925
www.mosaic-industries.com



Microcontrollers
- MCF51AC256
- MCF51CN128
- MCF51JM128
- MCF51QE128
- MC9S08AC128
- MC9S08JM64
- MC9S08QE128

Modules
- Bluetooth
- RS-232
- USB
- RF
- 2 x 20 LCD
- 4 x 4 Keypad
- Motor Control
- Real Time Clock
- Power Supply

GREAT For:
- Projects
- Breadboards
- Experiments
- Labs
- Prototypes

Kits Available

Learn about BASIC ON BOARD

ATRIA *Technologies Inc.*
www.AtriaTechnologies.com

www.microcontroladores.com



USB-PIC'School
Laboratory for PIC's:
- Debugger/Programmer
- 70 exercises collection
- USB Interface

175 €



UNIVERSAL TRAINER:

Professional laboratory for developing designs and teaching Electronics

140 €

Ingeniería de Microsistemas Programados S.L.

MSE MICROSYSTEMS ENGINEERING

Alda. Mazarredo, 47 -1 º
48009 - Bilbao (Spain)
info@microcontroladores.com
www.microcontroladores.com

microEngineering Labs, Inc.
www.melabs.com 888-316-1753

PICBASIC PRO™ Compiler



Bridges the gap between ease of use and professional level results. The simplicity of the language allows both hobbyist and engineer to master it quickly, but the engineer appreciates what's happening under the hood. This is a true compiler that produces fast, optimized machine code that's stable and dependable.

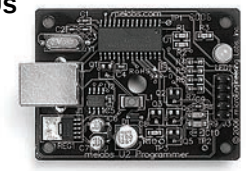
Supports:
More than 300 PIC® Microcontrollers
Direct Access to Internal Registers
In-Line Assembly Language
Interrupts in PICBASIC and Assembly
Built-In USB, I2C, RS-232 and More
Source Level Debugging in MPLAB

PICBASIC PRO™ Compiler $249.95

USB Programmer for PIC® MCUs

$89.95 *(as shown)*

RoHS Compliant



Programs PIC MCUs including low-voltage (3.3V) devices. Includes Software for Windows 98, Me, NT, XP, and Vista.

With Accessories for $119.95:
Includes Programmer, Software, USB Cable, and Programming Adapter for 8 to 40-pin DIP.

See our full range of products, including books, accessories, and components at:
www.melabs.com



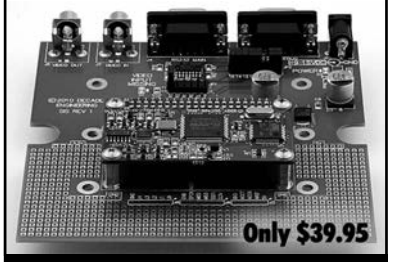
$795
introductory price

- ¼ " Ultrasonic Microphone & Amplifier
- 500 Hz – 100KHz +/- 3 dB BSL
- Sensitivity: 1 V/P

KarlZurn@gmail.com
321-613-3862

FLORIDA RESEARCH INSTRUMENTS, INC.
www.FloridaResearchInstruments.com

ALL ELECTRONICS CORPORATION
Electronic and Electro-mechanical Devices, Parts and Supplies.
Wall Transformers, Alarms, Fuses, Relays, Opto Electronics, Knobs, Video Accessories, Sirens, Solder Accessories, Motors, Heat Sinks, Terminal Strips, L.E.D.S., Displays, Fans, Solar Cells, Buzzers, Batteries, Magnets, Cameras, Panel Meters, Switches, Speakers, Peltier Devices, and much more....
www.allelectronics.com
Free 96 page catalog
1-800-826-5432

NEW SIS Board
Turbocharges the development of low-cost modular video text and graphics generators using BOB-4!


www.decadenet.com
DECADE ENGINEERING
503-743-3194 Turner, OR, USA

Disk On Chip
Ready for Delivery

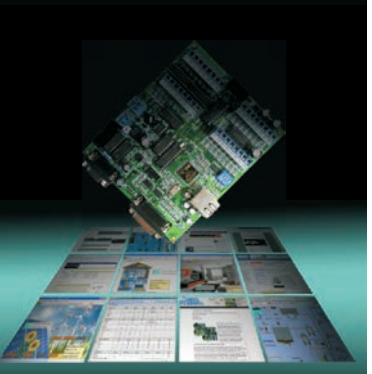

From 16MB to 128MB Available!
Call 530-297-6073 Sales@jkmicro.com
www.jkmicro.com
JK microsystems, Inc.
International Orders Welcome

IC Device & Package Converters
Quick-Turn Custom Solutions
• Upgrade Your PCB • Design and assembly - Fast
• Fix Design Problem
• Replace Obsolete Part
• All SMT Package Types
Economical Replacement for QFP, PLCC, BGA, SOIC IC's

RoHS
Ironwood ELECTRONICS 1-800-404-0204
www.ironwoodelectronics.com

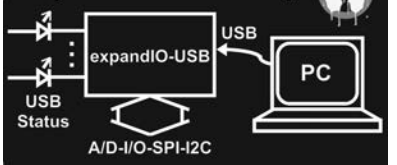


Add a Touch Screen to Your Embedded Product
• No special OS or Library Required.
• Programming GUI is Simple.
• Development Kit = Up and Running in Days.
Learn more at reachtech.com, call 510-770-1417 or email info@reachtech.com.
20 REACH TECHNOLOGY INC. 1988-2008
REACH TECHNOLOGY INC.
Buy vs. Build? Download Handbook at reachtech.com.

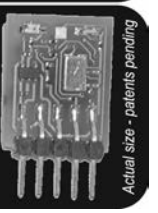
INTEGRATED
PLCs with Integrated Ethernet

Integrated Features :
- ETHERNET / Modbus TCP/IP
- 16 or 32 digital I/Os
- 10 analog I/Os
- RS232 and RS485
- LCD Display Port
- I/O Expansion Port
- Ladder + BASIC Programming
$229 and $295
before OEM Qty Discount
tel : 1 877 TRI-PLCS
web : www.tri-plc.com/cci.htm
TRI TRIANGLE RESEARCH INTERNATIONAL

CCS Optimized C Compiler for PIC® MCU
· Decrease code size by 40%
· Use less expensive MCU with smaller program memory
· Save hundreds of dollars in component costs in your projects

CCS Compilers Save You Money
Compilers Starting at $150
CCS www.ccsinfo.com/SAVE
sales@ccsinfo.com
262.522.6500 x35

Revolutionary new expandIO-USB chip


• Ideal for adding USB to sensors & peripherals
• No drivers needed for Windows, Mac, Linux
• No microcontroller programming required
• Also check out our USB-232 USB to UART
www.hexwax.com - Buy from Mouser & Farnell

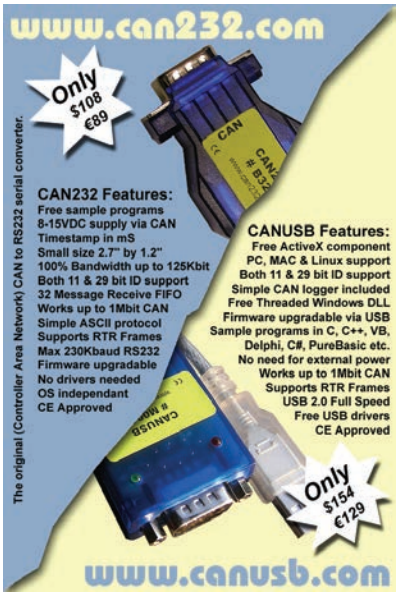
Amazing PIC programmer
Most devices supported
ICSP, SQTP, & copy limits
$32 at Digikey & Mouser
www.flexipanel.com


i2cStick™



Discover I²C
Read, Write, Connect...
MCC
Micro Computer Control
www.mcc-us.com

www.can232.com



Only $108 €89

CAN232 Features:
Free sample programs
8-15VDC supply via CAN
Timestamp in mS
Small size 2.7" by 1.2"
100% Bandwidth up to 125Kbit
Both 11 & 29 bit ID support
32 Message Receive FIFO
Works up to 1Mbit CAN
Simple ASCII protocol
Supports RTR Frames
Max 230Kbaud RS232
Firmware upgradable
No drivers needed
OS independant
CE Approved

CANUSB Features:
Free ActiveX component
PC, MAC & Linux support
Both 11 & 29 bit ID support
Simple CAN logger included
Free Threaded Windows DLL
Firmware upgradable via USB
Sample programs in C, C++, VB, Delphi, C#, PureBasic etc.
No need for external power
Works up to 1Mbit CAN
Supports RTR Frames
USB 2.0 Full Speed
Free USB drivers
CE Approved

The original (Controller Area Network) CAN to RS232 serial converter.

Only $154 €129

www.canusb.com

MaxSonar
Perfect for OEMs & Engineers



XL-MaxSonar-EZ
• High acoustic power
• Real-time calibration
• Tiny size
• $39.95 USD qty. 1

MaxSonar-WR (IP67)
• Industrial packaging
• Weather resistant
• Standard ¾" fitting
• $99.95 USD qty. 1





LV-MaxSonar-EZ
• Power-up calibration
• Smallest MaxSonar
• Low power, 2.5V-5.5V
• $29.95 USD qty. 1

MaxSonar-WRC (IP67)
• Compact packaging
• Weather resistant
• Quality narrow beam
• $99.95 USD qty. 1



www.maxbotix.com

FlashCore-B [FB]™
Add embedded mass data storage, 16-bit ADCs to your product

$79 Qty 1 $34 OEM



• 2.1" x 2.4", C/C++ programmable.
• CompactFlash interface with FAT file system.
• 16-bit ADCs, DACs, RS232 and TTL I/Os.
• Ultra-low quiescent current for battery power.

50+ Low Cost Controllers with ADC, DAC, 18 UARTs, 300 I/Os, solenoid, relays, CompactFlash, LCD, DSP motion control. Custom board design. Save time and money.

TERN INC. 1724 Picasso Ave., Suite A, Davis, CA 95616 USA
Tel: 530-758-0180 • Fax: 530-758-0181 www.tern.com • sales@tern.com

Embedded Ethernet
ONLY $98



NEW! picoFlash

• 10Base-T Ethernet
• 186 Processor @ 40 MHz
• DOS w/ Flash File System
• 16 Digital I/O • 5V DC
• Console / Debug Port
• Hardware Clock/Calendar
• Socket for DiskOnChip
• 512K DRAM & 512K Flash
• (2) Serial Ports
• (2) 16-bit Timers
• Watchdog Timer
• 3.75" x 2.50"

$129 Development System

Development Kit Includes:
• picoFlash Controller
• Borland C/C++ 4.52
• TCP/IP, PPP & Web Server
• Serial Driver Library
• AC Adapter and Cables

Call 530-297-6073 Email sales@jkmicro.com
On the web at www.jkmicro.com

JK microsystems

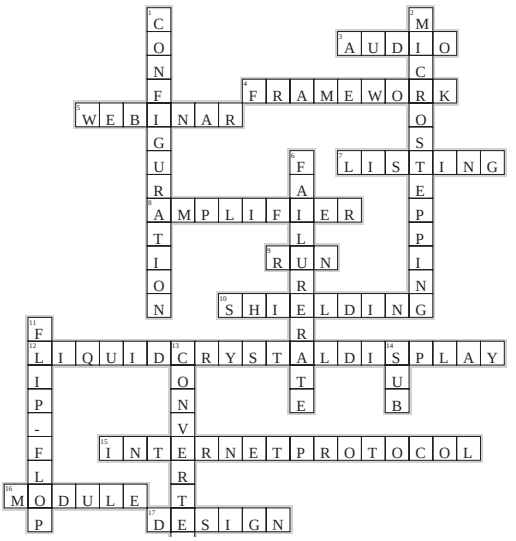
CROSSWORD ANSWERS from Issue 256

Across

3. **AUDIO**—Not visual
4. **FRAMEWORK**—The foundation of an application
5. **WEBINAR**—An online seminar
7. **LISTING**—A series of lines of code
8. **AMPLIFIER**—A device used to increase power
9. **RUN**—Fast movement, a sequence or course
10. **SHIELDING**—A barrier of conductive or magnetic matter
12. **LIQUIDCRYSTALDISPLAY**—LCD [three words]
15. **INTERNETPROTOCOL**—IP [two words]
16. **MODULE**—A unit of measurement
17. **DESIGN**—A plan or specification

Down

1. **CONFIGURATION**—An arrangement of units
2. **MICROSTEPPING**—Decreases vibration
6. **FAILURERATE**—The amount of times a component is unsuccessful [two words]
11. **FLIP-FLOP**—A circuit with two stable states
13. **CONVERTER**—Changes analog to digital and vice versa
14. **SUB**—Below



Elektor Print

Classical reading: on paper



+

Elektor Digital

Contemporary & innovative reading: on PC, notebook or tablet

=

Elektor PLUS

Comprehensive reading: at home and on the road

Read Elektor with the premium PLUS subscription!

Subscribe now or upgrade: www.elektor.com/usa

INDEX OF ADVERTISERS

Page		Page		Page		Page	
41	2012 International CES	78	Elektor	75	Ingenieria De Microsistemas	7	Mouser Electronics, Inc.
76	All Electronics Corp.	33	Elprotronic, Inc.	76	Ironwood Electronics	5	Newhaven Display International
37	AP Circuits	13	Elsevier, Inc.	23	Jeffrey Kerr, LLC	C2	NetBurner
61	ARM	37	EMAC, Inc.	9, 33	JK microsystems, Inc.	C4	Parallax, Inc.
75	Atria Technologies, Inc.	27	Embedded Developer	76, 77	JK microsystems, Inc.	18	Pico Technology Ltd.
C3	Audio Amateur	49	ExpressPCB	17	Labcenter Electronics	43	Pololu Corp.
65	Comfile Technology, Inc.	76	FlexiPanel Ltd.	77	Lawicel AB	76	Reach Technology, Inc.
76	Custom Computer Services, Inc.	75	Florida Reasearch Instruments, Inc.	77	Maxbotix, Inc.	63	Saelig Co., Inc.
76	Decade Engineering	35	GHI Electronics LLC	77	MCC, Micro Computer Control	45	SiliconRay
31	DesignSpark chipKIT™ Challenge	45	Grid Connect, Inc.	1	Microchip Technology, Inc.	2, 3	Technologic Systems
9	Earth Computer Technologies, Inc.	19	Humandata Ltd.	75	microEngineering Labs, Inc.	77	Tern, Inc.
10, 11	Elektor	57, 73	IC Bank	75	Mosaic Industries, Inc.	76	Triangle Research, Inc.

PREVIEW of January Issue 258

Theme: Embedded Applications

A Different Timeshare

mbos: A Real-Time Operating System for mbed

SMD Air Bath: A Surface-Mount Device for Temperature Control

NakedCPU (Part 1): A Roadmap for Navigating Hardware Documentation

---

THE CONSUMMATE ENGINEER **Building a Robust Power Supply (Part 1):** Power Dissipation

FROM THE BENCH **Quantifying Precipitation:** A Device to Measure Rainfall Volume

EMBEDDED UNVEILED **Digital X-ray Sensors**

LESSONS FROM THE TRENCHES **Analog-to-Digital Converters and Signal Filtering**

ATTENTION ADVERTISERS

February Issue 259 Deadlines

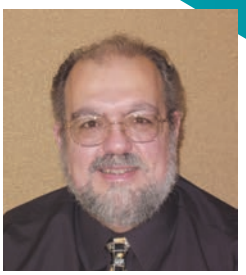
Space Close: Dec. 12
Materials Close: Dec. 19

Theme
Wireless Communications

Bonus Distribution
APEC; CTIA Wireless

Call Peter Wostrel
now to reserve your space!
800.454.3741 or 978.281.7708
e-mail: peter@smmarketing.us

PRIORITY INTERRUPT



by Steve Ciarcia, Founder and Editorial Director

40 Years and Still Going Strong

November 2011 marked the 40th anniversary of the public announcement of the 4004 microprocessor chip by Intel Corp. The chip that started it all had very modest specifications: a 4-bit data bus, up to 32,768 bits of program memory (4 KB in four chips), and up to 5,120 bits of read-write memory (1,280 4-bit nibbles in 16 chips). It required eight two-phase clock cycles to complete a single instruction cycle that took 10.8 microseconds. Five of the clock cycles were required just to fetch a byte of program memory. At 92.6K instructions per second, this snail speed seems laughable by today's standards, even for small processors. On the plus side, it also represented an entirely new computer packaging concept that kick-started a new industry called embedded computing.

Indeed, the 4004 was developed specifically as an embedded processor. It was initially designed as the brain inside a new line of "Busicom" desktop calculators being developed at the time by Nippon Calculating Machine Corp. (NCMC), a Japanese company. With perfect foresight, in 1971 Intel negotiated a new contract with NCMC that enabled Intel to market the chip to the public as a general-purpose microprocessor.

Engineers at other companies, myself included, began working with the 4004 and its successors—the 4040 and the 8008 (the latter introduced just five months later)—for all kinds of applications. After all, computer control wasn't an alien concept. By then, minicomputers from the likes of Digital Equipment Corp., Data General Corp., and Control Data Corp. had been used to control both laboratory and industrial processes for years. The microprocessor was just a better-packaged version of an existing concept that represented a huge opportunity to reduce the cost of those systems. Superior price/performance wins every time.

It didn't take long for these engineers to start tinkering with other "unofficial" projects involving microprocessors, and many soon realized they were powerful enough for general-purpose computing. Entrepreneurs among the group created the first self-contained desktop "microcomputer," complete with mass storage and a terminal for user interface. By 1976, Apple, Inc. was founded and, in 1980, even industry giant IBM had jumped onto the bandwagon with the IBM PC. There was no turning back once this powerful system tool hit the mass market. As soon as people discovered they could purchase an inexpensive system to do the things they had been doing on expensive mainframes and minicomputers, they couldn't get enough of the "microcomputer."

A bifurcation in the development of the microprocessor began around this time. Some chip and software developers continued focusing on the desktop applications (e.g., word processing, spreadsheets, and databases) while others redirected their attention to embedded applications. For a while, the same chips were used in both areas, but eventually hardware optimizations for one area created chips that weren't really suitable for the other area. In the desktop world, this led to the x86 architecture, which is still pervasive today. Generally speaking, x86 really only gets embedded in the form of a more-or-less complete "PC on a board."

In the embedded world, performance enhancements centered around real-time applications, including digital signal processing (DSP) of audio, pictures, and video streaming. The modern smartphone is the epitome of this, integrating the basic telephone functionality and access to the digital cellphone network with a digital still/video camera and a GPS receiver. The amount of DSP going on inside a smartphone is just astounding, not to mention the glitzy graphics-intensive user interfaces these devices incorporate.

However, the vast majority of embedded microprocessors exist in the form of the low- to mid-range single-chip microcomputers that handle pedestrian day-to-day tasks, such as managing your microwave oven, the engine and accessories in your car, the remote control for your TV set, and thousands of other similar applications. These chips range from tiny units in six-pin packages the size of a grain of rice to more substantial—yet still reasonably inexpensive—chips with 32-bit processors that have better throughput and more memory than the desktop systems of the 1980s and perhaps even the 1990s. ARM processors are a perfect example.

It's hard to think of any modern electronic device that hasn't had its design or applications "smartened up" along the way. If there is a message in this 40-year timeline, it's that a microprocessor that initially made a great calculator proved to have a whole lot more potential when exploited and enhanced by others. We'd like to think that for the last 30-plus years, *Circuit Cellar* has been a prime vehicle for bringing attention to technologies worth exploiting and enhancing. We pledge to keep up the good work for the next 40 years.

steve.ciarcia@circuitcellar.com



Subscribe to *audioXpress* magazine!

Do your **electronics speak** to you? Are the words **"audio", "vacuum tubes",** and **"speaker technology"** music to your ears?

Then you should be reading *audioXpress!*

Recently acquired by The Elektor Group, *audioXpress* has been providing engineers with incredible audio insight, inspiration and design ideas for over a decade. If you're an audio enthusiast who enjoys speaker building and amp design, or if you're interested in learning about tubes, driver testing, and vintage audio, then *audioXpress* is the magazine for you!

What will you find in *audioXpress*?

- In-depth interviews with audio industry luminaries
- Recurring columns by top experts on speaker building, driver testing, and amp construction
- Accessible engineering articles presenting inventive, real-world audio electronics applications and projects
- Thorough and honest reviews about products that will bring your audio experiences to new levels

Choose from print delivery, digital, or a combination of both for maximum accessibility.

Subscribe to *audioXpress* at
www.cc-webshop.com
today!



audioXPRESS

2011 Gift Guide



For all your holiday gift giving (and gift getting!) needs, here is a tailored group from our vast selection of microcontrollers, robots, kits, sensors and accessories. For complete product catalog, visit www.parallax.com *(Clockwise from upper left)*

Boe-Bot Robot Kit (USB Version; #28832; $159.99) Learn programming and robotics with this BASIC Stamp-controlled robot kit. Serial with USB adapter version also available (#28132)

PING))) Ultrasonic Sensor (#28015; $29.99) Easy method of distance measurement for robotics, security, and infrared replacement applications.

S2 (Scribbler 2) Robot (#28136; $129.99) The S2 is pre-programmed with 8 demo modesand has a colorful GUI programming interface for your PC.

VPN1513 GPS Receiver module (#28506; $59.99) An open source, customizable GPS solution that can track up to 20 satellites.

Memsic 2125 Dual-axis Accelerometer (#28017; $29.99) A dual-axis thermal accelerometer to measure tilt, acceleration, rotation, and vibration; range of ±3 g.

BASIC Stamp Activity Kit (#90005; $79.99) Learn circuit-building and programming at the same time with the BASIC Stamp microcontroller!

PIR Sensor (#555-28027; $10.99) A passive infrared motion sensor for your microcontroller project.

RFID Reader and Tag Sampler Kit (Serial; #32390; $42.99) Easily read passive RFID transponder tags.

Order online anytime at **www.parallax.com** or call toll-free at 888-512-1024 (M-F, 7AM-5PM, PDT).

Friendly microcontrollers, legendary resources.™

Prices subject to change without notice. BASIC Stamp, and Boe-Bot are trademarks of Parallax Inc. Parallax and the Parallax logo are trademarks of Parallax Inc.

PARALLAX INC

www.parallax.com

"ParallaxInc" on Twitter, Facebook, and YouTube