




# circuit cellar


# DATA-GATHERING DESIGNS

Wearable medical computing,  
thermal monitoring, and more




- Hackerspaces | Q&A: Biochip Design, Smart Homes, and More
- Solar Power Setup | USB Switching Power Supply | Thermal Monitoring
- Mobile Device Power Management | Strobe Photography Techniques |
- Sync Atomic Clocks | LCD-Enhancing Peripherals ■ Mobile Health Sensors

# Ethernet Core Modules with High-Performance Connectivity Options



- **MOD5270**  
147.5 MHz processor with 512KB Flash & 8MB RAM · 47 GPIO  
3 UARTs · I<sup>2</sup>C · SPI
- **MOD5234**  
147.5 MHz processor with 2MB flash & 8MB RAM · 49 GPIO · 3 UARTs  
I<sup>2</sup>C · SPI · CAN · eTPU (for I/O handling, serial communications,  
motor/timing/engine control applications)
- **MOD54415**  
250 MHz processor with 32MB flash & 64MB RAM · 42 GPIO · 8 UARTs  
5 I<sup>2</sup>C · 3 SPI · 2 CAN · SSI · 8 ADC · 2 DAC · 8 PWM · 1-Wire<sup>®</sup> interface
- **NANO54415**  
250 MHz processor with 8MB flash & 64MB RAM · 30 GPIO · 8 UARTs  
4 I<sup>2</sup>C · 3 SPI · 2 CAN · SSI · 6 ADC · 2 DAC · 8 PWM · 1-Wire<sup>®</sup> interface

**Add Ethernet connectivity to an existing product, or use it as your product's core processor**



**The goal:** Control, configure, or monitor a device using Ethernet

**The method:** Create and deploy applications from your Mac or Windows PC. Get hands-on familiarity with the NetBurner platform by studying, building, and modifying source code examples.

**The result:** Access device from the Internet or a local area network (LAN)

**The NetBurner Ethernet Core Module** is a device containing everything needed for design engineers to add network control and to monitor a company's communications assets. For a very low price point, this module solves the problem of network-enabling devices with 10/100 Ethernet, including those requiring digital, analog and serial control.

MOD5270-100IR.....\$69 (qty. 100)	NNDK-MOD5270LC-KIT.....\$99
MOD5234-100IR.....\$99 (qty. 100)	NNDK-MOD5234LC-KIT.....\$249
MOD54415-100IR.....\$89 (qty. 100)	NNDK-MOD54415LC-KIT.....\$129
NANO54415-200IR....\$69 (qty. 100)	NNDK-NANO54415-KIT.....\$99

**NetBurner Development Kits** are available to customize any aspect of operation including web pages, data filtering, or custom network applications. The kits include all the hardware and software you need to build your embedded application.

➤ **For additional information please visit**  
<http://www.netburner.com/kits>



# TFT Displays

Available in 1.8"-7.0" sizes

- Capacitive and Resistive touch options
- MVA, IPS, and TN technologies
- WVGA, QVGA, and VGA resolutions
- Standard stocked displays
- Development kits available
- Competitive pricing with low MOQs
- Custom options available
- Superior engineering support



NEWHAVEN DISPLAY INTERNATIONAL

For more information visit: [circuitcellar.newhavendisplay.com](http://circuitcellar.newhavendisplay.com)

Contact us at: 847-844-8795



RoHS  
Compliant

Issue 288 July 2014 | ISSN 1528-0608

CIRCUIT CELLAR® (ISSN 1528-0608) is published monthly by:

Circuit Cellar, Inc.  
111 Founders Plaza, Suite 300  
East Hartford, CT 06108

Periodical rates paid at East Hartford, CT, and additional offices.

One-year (12 issues) subscription rate US and possessions \$50, Canada \$65, Foreign/ ROW \$75. All subscription orders payable in US funds only via Visa, MasterCard, international postal money order, or check drawn on US bank.

#### SUBSCRIPTIONS

Circuit Cellar, P.O. Box 462256, Escondido, CA 92046

**E-mail:** circuitcellar@pcspublink.com

**Phone:** 800.269.6301

**Internet:** circuitcellar.com

**Address Changes/Problems:** circuitcellar@pcspublink.com

**Postmaster:** Send address changes to  
Circuit Cellar, P.O. Box 462256, Escondido, CA 92046

#### ADVERTISING

Strategic Media Marketing, Inc.  
2 Main Street, Gloucester, MA 01930 USA

**Phone:** 978.281.7708

**Fax:** 978.281.7706

**E-mail:** circuitcellar@smmarketing.us  
Advertising rates and terms available on request.

#### New Products:

New Products, Circuit Cellar, 111 Founders Plaza, Suite 300  
East Hartford, CT 06108, E-mail: newproducts@circuitcellar.com

#### HEAD OFFICE

Circuit Cellar, Inc. 111 Founders Plaza, Suite 300  
East Hartford, CT 06108  
Phone: 860.289.0800

#### COVER PHOTOGRAPHY

Chris Rakoczy, www.rakoczyphoto.com

#### COPYRIGHT NOTICE

Entire contents copyright © 2014 by Circuit Cellar, Inc. 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.

#### DISCLAIMER

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 infringing liability for constructing or operating such devices.

© Circuit Cellar 2014 Printed in the United States

## ADVANCES IN WEARABLE ELECTRONICS

The concept of the wireless body-area network (WBAN), a network of wireless wearable computing devices, holds great promise in health-care applications. Such a network could integrate implanted or wearable sensors that provide continuous mobile health (mHealth) monitoring of a person's most important "vitals"—from calorie intake to step count, insulin to oxygen levels, and heart rate to blood pressure. It could also provide real-time updates to medical records through the Internet and alert rescue or health-care workers to emergencies such as heart failures or seizures.

Conceivably, the WBAN would need some sort of controller, a wearable computational "hub" that would track the data being collected by all the sensors, limit and authorize access to that information, and securely transmit it to other devices or medical providers.

This issue features an essay by Clemson University researcher Vivian Genaro Motti, who discusses her participation in the federally funded Amulet project (p. 80). Amulet's Clemson and Dartmouth College research team is prototyping pieces of "computational jewelry" that can serve as a body-area network's mHealth hub while being discreetly worn as a bracelet or pendant. Motti's essay elaborates on Amulet's hardware and software architecture.

Motti isn't the only one aware of the keen interest in WBANs and mHealth. In an interview on p. 8, Shiyan Hu, a professor whose expertise includes very-large-scale integration (VLSI), says that many of his students are exploring "portable or wearable electronics targeting health-care applications."

Today's mHealth market is evident in the variety of health and fitness apps available for your smartphone. But the most sophisticated mHealth technologies are not yet accessible to embedded electronics enthusiasts. (However, Amulet has created a developer prototype with an easily accessible board, pictured below.)

Market demand tends to increase access to new technologies. A BCC Research report predicts the mHealth market, which hit \$1.5 billion in 2012, will increase to \$21.5 billion by 2018. Evolving smartphones, better wireless coverage, and demands for remote patient monitoring are fueling the growth. So you can anticipate more designers and developers will be exploring this area of wearable electronics.



**Mary Wilson**

mwilson@circuitcellar.com

## THE TEAM

#### EDITOR-IN-CHIEF

C. J. Abate

#### MANAGING EDITOR

Mary Wilson

#### ASSOCIATE EDITOR

Nan Price

#### ART DIRECTOR

KC Prescott

#### ADVERTISING COORDINATOR

Kim Hopkins

#### PRESIDENT

Hugo Van haecke

#### PUBLISHER

Don Akkermans

#### ASSOCIATE PUBLISHER

Shannon Barraclough

#### COLUMNISTS

Jeff Bachiochi, Ayse K.

Coskun, Bob Japenga, Robert Lacoste, Ed Nisley, George Novacek, Colin O'Flynn

#### FOUNDER

Steve Ciarcia

#### PROJECT EDITORS

Chris Coulston,  
Ken Davidson,  
David Tweed

#### CUSTOMER SERVICE

Debbie Lavoie

**US/UK**

Don Akkermans  
+31 46 4389444  
d.akkermans@elektor.com

**ELEKTOR LABS**

Wisse Hettinga  
+31 46 4389428  
w.hettinga@elektor.com

**GERMANY**

Ferdinand te Walvaart  
+49 241 88 909-17  
f.tewalvaart@elektor.de

**FRANCE**

Denis Meyer  
+31 46 4389435  
d.meyer@elektor.fr

**NETHERLANDS**

Harry Baggen  
+31 46 4389429  
h.baggen@elektor.nl

**SPAIN**

Eduardo Corral  
+34 91 101 93 95  
e.corral@elektor.es

**ITALY**

Maurizio del Corso  
+39 2.66504755  
m.delcorso@inware.it

**SWEDEN**

Wisse Hettinga  
+31 46 4389428  
w.hettinga@elektor.com

**BRAZIL**

João Martins  
+31 46 4389444  
j.martins@elektor.com

**PORTUGAL**

João Martins  
+31 46 4389444  
j.martins@elektor.com

**INDIA**

Sunil D. Malekar  
+91 9833168815  
ts@elektor.in

**RUSSIA**

Nataliya Melnikova  
+7 965 395 33 36  
elektor.russia@gmail.com

**TURKEY**

Zeynep Köksal  
+90 532 277 48 26  
zkoksal@beti.com.tr

**SOUTH AFRICA**

Johan Dijk  
+31 6 1589 4245  
j.dijk@elektor.com

**CHINA**

Cees Baay  
+86 21 6445 2811  
ceesbaay@gmail.com

**OUR NETWORK**

VOICE COIL

audioexpress

elektor

**SUPPORTING COMPANIES**

Accutrace	7	HuMANDATA, Ltd.	78
Advanced Micro Circuits	79	IAR Systems	15
All Electronics Corp.	78	Imagineering, Inc.	C4
AP Circuits	43	Ironwood Electronics	78
Beta LAYOUT, Ltd.	51	MaxBotix, Inc.	79
CadSoft Computer GmbH	17	microEngineering Labs, Inc.	79
Cleverscope	51	MyRO Electronic Control Devices, Inc.	78
CTIA Wireless	31	NetBurner, Inc.	C2
Custom Computer Services	78	Newhaven Display International	1
Elektor	22, 23	R.E. Smith, Inc.	47
Elektor	63	RS Components	28-30
EMAC, Inc.	19	Saelig Co., Inc.	19
ExpressPCB	37	Technologic Systems	9, 11
Front Panel Express	43	TechTools	79
GHI Electronics, LLC	61	Triangle Research International, Inc.	79

**NOT A SUPPORTING COMPANY YET?**


Contact Peter Wostrel (circuitcellar@smmarketing.us, Phone 978.281.7708, Fax 978.281.7706) to reserve your own space for the next edition of our members' magazine.

## CONTENTS



JULY 2014 • ISSUE 288

## INTERNET & CONNECTIVITY



ANT+ WIRELESS TRANSMISSION COORDINATOR PROTOTYPE

### CC COMMUNITY

06 : CC WORLD

## 08 : QUESTIONS &amp; ANSWERS

Teaching, CAD Research, and VLSI Innovation

By Nan Price

Researcher Shiyun Hu on very-large-scale integration (VLSI), smart home monitoring, and biochip design

### INDUSTRY & ENTERPRISE

13 : PRODUCT NEWS

## 19 : CLIENT PROFILE

TechTools (Rowlett, TX)

### FEATURES

## 20 : Personal Solar Power Setup (Part 1)

Solar Panels, Batteries, and Storage

By Steve Hendrix

Install solar panels and integrate reliable backup power

## 32 : Automotive Isolated Switching Power Supply

By Tom Struzik

SPICE simulation is a key to designing an in-car, high-current USB power supply for smartphones and more

## 40 : Build a Thermal Monitoring Network

By Carlo Tauraso

Use ANT+ wireless technology to implement a real-time temperature monitoring system

### COLUMNS

## 48 : GREEN COMPUTING

Power Management in Mobile Devices

By Ayse K. Coskun


Learn about efficient power control mechanisms

## 54 : THE CONSUMMATE ENGINEER

WWVB Clock Revisited


By George Novacek

Build a GPS clock standard and WWVB signal monitor to synchronize radio-controlled atomic clocks



SENSOR-BASED SMART VIDEO MONITORING SYSTEM


# CONTENTS



GPS CLOCK STANDARD WITH AN XBEE TRANSCEIVER



LCD WITH A "COG" DISPLAY AND A BACKLIGHTING MODULE



BREADBOARD LAYOUT FOR A STROBE FLASH UNIT

## 58 : ABOVE THE GROUND PLANE

### Strobe Photography

By Ed Nisley

An Arduino board and MOSFET help trigger a camera and time flashes to capture a rapidly moving object


## 66 : FROM THE BENCH

### LCD User Interface

Reconfiguring and Serial Backpacking

By Jeff Bachiochi

Serial backpacking can make LCDs more user-friendly



TEST BOARD FOR AN ISOLATED SWITCHING POWER SUPPLY

## TESTS & CHALLENGES

74 : CROSSWORD

75 : TEST YOUR EQ



AMULET DEVELOPER PROTOTYPE AND REPLACEABLE BOARD

## TECH THE FUTURE

### 80 : Wearable Medical Computing and the Amulet Project

By Vivian Genaro Motti

A wearable, bracelet-style data "hub" for mobile health monitoring

## CC WORLD



# EIM RELAUNCHES AUDIOXPRESS.COM

By João Martins

In addition to *Elektor* magazine, Elektor International Media (EIM) publishes *Circuit Cellar* magazine and two monthly audio-related titles: *audioXpress* and *Voice Coil*. The latter is a periodical for the loudspeaker industry. *audioXpress* features content for audio enthusiasts and professionals alike: innovations in audio, audio electronics, and the best in DIY audio.

EIM's 2013 redesign, refocus, and relaunch of *audioXpress* was the first step in a larger plan to bring the publication into the digital era. The ultimate goal was to develop a complete web platform for the audio community that would reflect a range of activities and also launch online versions of EIM's

*Loudspeaker Industry Sourcebook (LIS)* and *World Tube Directory* publications.

For the new *audioXpress.com* website, our plan was to incorporate industry and product news with blog posts, interviews, magazine articles, project articles, reviews, and theory articles. As is typically the case with such initiatives, the project rapidly grew in size and scope, and we ultimately had to design a new platform from scratch to accommodate all our ambitions and team requests.

On Wednesday, May 21, we finally decided it was time to make *audioXpress.com* live, even though there was still work to do and we hadn't checked off every item on our wish list. EIM has been working hard to complete the project.

First and foremost, the new online platform is dynamic and ideally formatted to be viewed on any device, at different resolutions and screen sizes. Second, it connects the service level for our members with our dedicated team, while linking with social media services at every imaginable level. Finally, the new *audioXpress.com* platform addresses the market on a global scale and the information needs of the audio community in general, by serving an unlimited audience.

We invite *Circuit Cellar* and *Elektor* readers interested in all things audio to discover *audioXpress*, read about new audio innovations, learn about the latest technologies in audio electronics, and find interesting articles about DIY audio.



Visit the new *audioXpress.com* website for audio industry news, hi-fi product updates, and a variety of DIY projects.



The new *audioXpress.com* website features hi-fi product and industry news, DIY audio projects, interviews, in-depth articles about audio technology, and much more.



# ELECTRONICS HACKERSPACE PROFILES

By CC Staff

Hackerspaces and engineering clubs are popping up all over the world. And so *Circuit Cellar* is creating an online list of hackerspaces and including information about their locations, memberships, and equipment. Want to share information about your hackerspace, club, or group? Contact our editors if you'd like us to consider featuring your hackerspace on our website ([circuitcellar.com/hackerspaces-clubs](http://circuitcellar.com/hackerspaces-clubs)).



HackRVA is a 2,000-ft<sup>2</sup> space in Richmond, VA, that includes a meeting area, tech lab, and fab lab. (Source: [Hackrva.org](http://Hackrva.org))

The two most recent hackerspaces on the site were HackRVA and Denhac. The groups have 65 and 46 members, respectively.

HackRVA ([hackrva.org](http://hackrva.org)) is a Richmond, VA-based

hackerspace with two 3-D printers, a laser cutter, several soldering stations, oscilloscopes, a computer lab, and various power tools. "For the past three years, we've designed those PCB badges for the richSEC security conference," noted HackRVA member Aaron Nipper.

Denhac ([denhac.org](http://denhac.org)) is a Denver, CO-based space whose mission is "to create and sustain a local, community-driven, shared space that enables education, experimentation, and collaboration by applying the spirit of DIY to science, technology, engineering, and art." In its 2,500-ft<sup>2</sup> space, the group houses a small OpenStack-driven data center, a Cisco networking station, a 3-D printer station, an Internet radio station, and an electronics workstation with a variety of handy testing equipment.



*Circuit Cellar* is compiling an online list of electronics-focused hackerspaces and clubs. Let us know what you think.



PRINTED CIRCUIT BOARDS

THINK YOU CAN FIND PCB PRICES THAT BEAT OURS?

**WE DARE YOU.**



OUR GUARANTEE:

We are so confident in our PCB pricing, we dare you to find lower prices! If you do, we will match the price AND give you \$100 towards your next order!

*Visit us at [www.PCB4u.com](http://www.PCB4u.com) and see why our pricing can not be beaten!*

- From same day quick turn prototype to production in under 10 days
- Full CAD and CAM review plus design rule check on ALL Gerber files
- Materials: Fr4, Rigid, Flex, Metal Core (Aluminum), Polyimide, Rogers, Isola, etc.
- HDI Capabilities: Blind/Buried Microvias, 10+N+10, Via-in-Pad Technology, Sequential Lamination, Any Layer, etc.
- Our HDI Advantage: Direct Laser Drilling, Plasma De-Smear Technology, Laser Micro via Conductive Plate Shut, etc.

**Accutrace<sup>®</sup> inc.**

[www.PCB4u.com](http://www.PCB4u.com) • [sales@PCB4u.com](mailto:sales@PCB4u.com) • (408) 748-9600



## QUESTIONS & ANSWERS



# Teaching, CAD Research, and VLSI Innovation

## An Interview with Shiyan Hu

Shiyan Hu is an assistant professor in the Department of Electrical and Computer Engineering at Michigan Technological University. We discussed his research in the fields of computer-aided design (CAD), very-large-scale integration (VLSI), smart home monitoring, and biochip design.—Nan Price, Associate Editor

**NAN:** How long have you been at Michigan Technological University? What courses do you currently teach and what do you enjoy most about instructing?

**SHIYAN:** I have been with Michigan Tech for six years as an assistant professor. Effective September 2014, I will be an associate professor.

I have recently taught the graduate-level “Advanced Embedded System for Smart Infrastructure,” the graduate-level “Advanced Very-Large-Scale Integration (VLSI) Computer-Aided Design (CAD),” and the undergraduate-level “VLSI Design” courses.

The most exciting part about teaching is the interactions with students. For example, questions from students—although sometimes challenging—can be intriguing and it is fun to observe diversified thoughts. In addition, students taking graduate-level courses need to discuss their course projects with me. During the discussions, I can clearly see how excited they feel about their progress, which makes the teaching very enjoyable.

**NAN:** What “hot topics” currently interest your students?

**SHIYAN:** Students are very interested in embedded system designs for smart homes, including FPGA design and embedded programming for the scheduling of various smart home appliances to improve convenience and reduce the cost of electricity bills. I also frequently have meetings with students who are interested in portable or wearable electronics targeting health-care applications.

**NAN:** Describe your role as director of Michigan Tech’s VLSI CAD research lab.

**SHIYAN:** I have been advising a team of PhD

and MS students who conduct research in the area of VLSI CAD in the Electrical and Computer Engineering (ECE) department. A main research focus of our lab is VLSI physical design including buffer insertion, layer assignment, routing, gate sizing, and so forth. In addition, we have developed some embedded system prototypes such as sensor-based video monitoring and a 3-D mouse (see **Photo 1**).

There is also growing collaboration between our lab and the power system lab on the research of a CAD technique for smart-grid systems. The collaboration has led to an innovative optimization technique for an automatic feeder remote terminal unit that addresses cybersecurity attacks to smart power distribution networks. Further, there is an ongoing joint project on an FPGA-based embedded system for power quality prediction.

Although most of my time as the research lab director is spent on student mentoring and project management, our lab also contributes considerably to education in our department. For example, instructional and lab materials for the undergraduate “VLSI Design” course are produced by our lab.

**NAN:** Tell us more about your smart home research and the technique you developed to address cybersecurity problems.

**SHIYAN:** My smart home research emphasizes embedded systems that handle scheduling and cybersecurity issues. **Figure 1** shows a typical smart home system, which consists of various components such as household appliances, energy storage, photovoltaic (PV) arrays, and a plug-in hybrid electrical vehicle (PHEV) charger. Smart meters are installed at the customer side and connected to the smart power distribution system.

The smart meter can periodically

# SUPERIOR **EMBEDDED** SOLUTIONS



**DESIGN YOUR SOLUTION TODAY**  
**CALL 480-837-5200**


[www.embeddedARM.com](http://www.embeddedARM.com)

**NEW!**

## Single Board Computers

TS-7400-V2 Low Power & Low Cost

Backwards Compatible with TS-7400



pricing starts at  
**\$119**  
qty 1  
**\$84**  
qty 100

### Features:

- 450MHz ARM CPU
- Up to 256MB RAM
- 2GB NAND Flash
- 2x SD Card Socket
- 1x 10/100 Ethernet
- 2x USB Host
- 39x DIO, 4x ADC
- 4x Serial ports
- 1x BBRTC, 1x CAN
- 1x Temp Sensor

### Benefits:


- 2x faster, more features, reduced cost
- Low power with 120mW standby mode
- -40 to +85C, 100% soldered-on components
- Easy development w/ Debian and Linux 2.6
- Boots quickly to your Embedded Application
- Guaranteed available until 2025

**NEW!**

## Single Board Computers

TS-7250-V2 High Performance

Backwards Compatible with TS-72xx



shown w/  
optional  
micro  
SD Card

pricing starts at  
**\$199**  
qty 1  
**\$165**  
qty 100

### Features:

- Up to 1GHz ARM CPU
- 512MB RAM
- 4GB eMMC Flash
- 2x SD Card Socket
- 2x 10/100 Ethernet
- 2x USB Host
- 1x USB Device
- 6x Serial ports
- 75x DIO, 1x CAN
- 1x PC/104 Connector

### Benefits:

- Hardware Flexibility with On-board FPGA
- Several control I/O interfaces
- Launches your application in under a second
- Easy development w/ Debian and Linux 2.6
- High Data Reliability with DoubleStore
- -40 to +85C Industrial temperature range



We've never discontinued a product in 30 years



Embedded systems that are built to endure

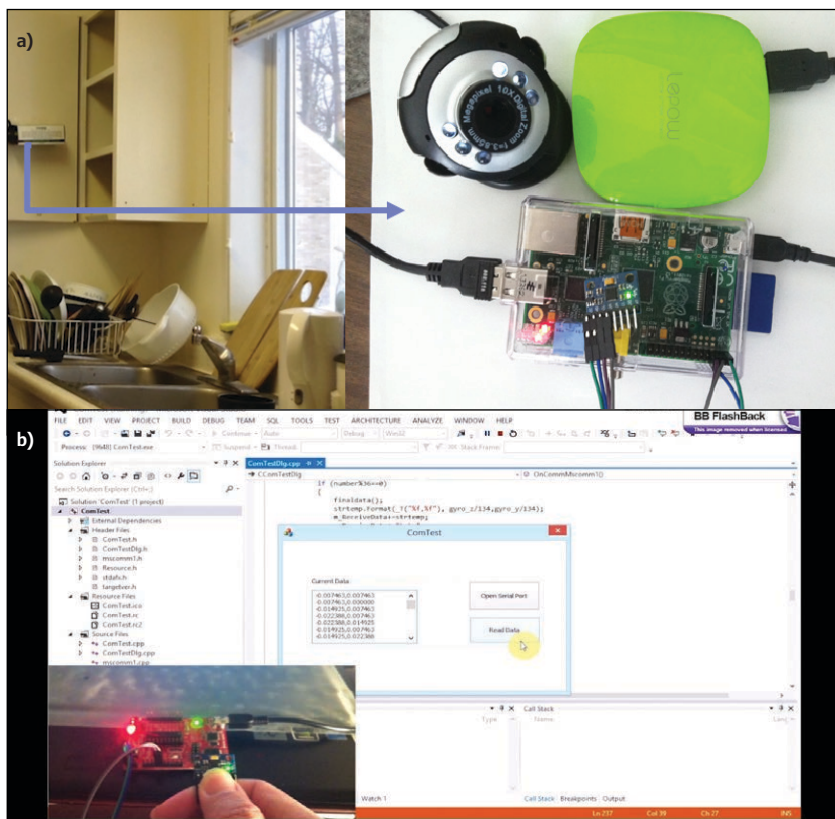


Support every step of the way with open source vision



Unique embedded solutions add value for our customers

## QUESTIONS & ANSWERS




**PHOTO 1**

Shiyan and a team of students he advised developed this sensor-based smart video monitoring system (a) and a 3-D mouse (b).

**FIGURE 1**

This smart home system architecture includes HVAC and several home appliances.



the undesired scheduling solutions. Cyber attackers can take advantage of this by scheduling their own energy-consuming tasks during the inexpensive hours, which would be expensive without a cyber attack. This is an interesting topic I am working on.

**NAN:** Describe your VSLI research.

**SHIYAN:** Modern ICs and chips are ubiquitous. Their applications include smartphones, modern home appliances, PCs, and laptops, as well as the powerful servers for big data storage and processing. In VLSI and system-on-a-chip (SoC) design, the layout design (a.k.a., physical design) often involves billions of transistors and is therefore enormously complex. Handling such a complex problem requires high-performance software automation tools (i.e., physical design automation tools) to achieve design objectives within a reasonable time frame. VLSI physical design is a key part of my research area.

**NAN:** Are you involved in any other areas of research?

**SHIYAN:** I also work on microfluidic biochip design. The traditional clinical diagnosis procedure includes collecting blood from patients and then sending it to laboratories, which require space and are labor-intensive and expensive, yet sometimes inaccurate.

The invention of the lab on a chip (a.k.a., biochip) technology offers some relief. The expensive laboratory procedures can be simply performed within a small chip, which provides much higher sensitivity and detection accuracy in blood sample analysis and disease diagnosis. Some point-of-care versions of these have already become popular in the market.

A major weakness of the prevailing biochip technology is that such a chip often has very limited functionality in terms of the quantities it can measure. The reason is that currently only up to thousands of biochemical reactions can be handled in a single biochip. Since the prevailing biochips are always manually designed, this seems to be the best one can achieve. If a single biochip could simultaneously execute a few biological assays corresponding to related diseases, then the clinical diagnosis would be much less expensive and more convenient to conduct. This is also the case when utilizing biochips for biochemical research and drug discovery.

My aim for this biochip research project is to largely improve the integration complexity of miniaturized components in a biochip

# SUPERIOR **EMBEDDED** SOLUTIONS




**DESIGN YOUR SOLUTION TODAY**  
**CALL 480-837-5200**

[www.embeddedARM.com](http://www.embeddedARM.com)

## Computer-on-Modules

State of the Art Embedded Design



pricing  
starts at  
**\$139**  
qty 1  
**\$99**  
qty 100

- Up to 1.2GHz ARM w/ 2GB RAM
- Fanless, low power designs
- All parts soldered on, no moving parts
- DoubleStore SD for reliable storage
- User-programmable FPGAs
- -40 to +85C Industrial temperature range
- Easy development w/ Android, Debian and Linux
- Sub-second boot to Linux
- COTS carrier boards available or design a custom solution with reduced design time and complexity

### Computer-on-Modules include:

- TS-4200: Atmel ARM9 w/ super low power
- TS-4600: 450MHZ low cost w/ 2 Ethernets
- TS-4710: Up to 1066MHZ PXA168 w/ video
- TS-4712: like TS-4710 + 2 Ethernets
- TS-4720: like TS-4712 + 4GB eMMC Flash
- TS-4800: 800MHz FreeScale iMX515 w/ video
- TS-4900: 1.2GHz QuadCore i.MX6 w/ WiFi

NEW!

NEW!

NEW!


NEW!

## Touch Panel Computers

Panel Mount or Fully Enclosed



series  
starts at  
**\$369**  
qty 100  
**\$409**  
qty 1



### Features can include:

- 5-inch, 7-inch and 10-inch touchscreens
- Fanless operation from -20 to +70C
- Up to 1.2Ghz ARM CPU
- Up to 2GB RAM, 4GB eMMC Flash
- 2x microSD with DoubleStore
- 2x Ethernet, 2x USB Host
- CAN, RS-232, SPI, I2C, DIO
- 1x RS485 2W-Modbus
- Optional cellular, WIFI & XBEE radios
- Support Android and Linux w/ sub-second boot
- Headphone connector & speaker



We've never discontinued a product in 30 years



Embedded systems that are built to endure



Support every step of the way with open source vision



Unique embedded solutions add value for our customers

## QUESTIONS & ANSWERS

to provide many more functionalities. The growing design complexity has mandated a shift from the manual design process toward a CAD process.

Basically, in the microfluidic biochip CAD methodology, those miniaturized components, which correspond to fundamental biochemical operations (e.g., mix and split), are automatically placed and routed using computer algorithms. This methodology targets minimizing the overall completion time of all biochemical operations, limiting the sizes of biochips, and improving the yield in the biochip fabrication. In fact, some results from our work were recently featured on the front cover of *IEEE Transactions on Nanobioscience* (Volume 13, No. 1, 2014), a premier nanobioscience and engineering journal. In the future, we will consider inserting on-chip optical sensors to provide real-time monitoring of the biological assay execution, finding possible errors during execution, and dynamically reconfiguring the biochip for error recovery.

**NAN: You've earned several distinctions and awards over the last few years. How do these acknowledgments help your research?**

**SHIYAN:** Those awards and funding certainly help me a lot in pursuing the research of fascinating topics. For example, I am a 2014 recipient of the NSF CAREER award, which is widely regarded as one of the most prestigious honors for up-and-coming researchers in science and engineering.

My five-year NSF CAREER project will focus on carbon nanotube interconnect-based circuit design. In the prevailing 22-nm technology node, wires are made from coppers and such a thin copper wire has a very small cross-section area. This results in large wire resistance and large interconnect delay. In fact, the interconnect delay has become the limiting factor for chip timing. Due to the fundamental physical limits of copper wires, novel on-chip interconnect materials (e.g., carbon nanotubes and graphene nanoribbons) are more desirable due to their many salient features (e.g., superior conductivity and resilience to electromigration).

To judiciously integrate the benefits from both nanotechnology interconnects and copper interconnects, my NSF CAREER project will develop a groundbreaking physical layout codesign methodology for next-generation ICs. It will also develop various physical design automation techniques as well as a variation-aware codesign technique for the new methodology. This project aims to integrate the pioneering nanotechnologies into the practical circuit design and it has the potential to contribute to revolutionizing the prevailing

circuit design paradigm.

**NAN: Give us some background information. When did you first become interested in computer engineering?**

**SHIYAN:** I started to work on computer engineering when I entered Texas A&M University conducting research with professor Jiang Hu, a leading expert in the field of VLSI physical design. I learned a lot about VLSI CAD from him and I did several interesting research projects including buffer insertion, gate sizing, design for manufacturability, and post silicon tuning. Through his introduction, I also got the chance to collaborate with leading experts from IBM Research on an important project called "slew buffering."


**NAN: Tell us more about your work at IBM Research.**

**SHIYAN:** As VLSI technology scales beyond the 32-nm node, more devices can fit onto a chip, which implies continued growth of design size. The increased wire delay dominance due to finer wire widths makes design closure an increasingly challenging problem.

Buffer insertion, which improves an IC's timing performance by inserting non-inverting buffers or inverting buffers (a.k.a., inverters), has proven to be indispensable in interconnect optimization. It has been well documented that typically more than 25% of gates are buffers in IBM ASIC designs.

Together with my collaborators at IBM Research, I proposed a new slew buffering-driven dynamic programming technique. The testing with IBM ASIC designs demonstrated that our technique achieves a more than 100x speed increase compared to the classical buffering technique while still saving buffers. Therefore, the slew buffering-driven technique has been implemented and deployed into the IBM physical design flow as a default option.

IBM researchers have witnessed that the slew buffering technique contributes to a great reduction in the turnaround time of the physical synthesis flow. In addition, more extensive deployment of buffering techniques leads to superior design quality. Such an extensive buffer deployment-based interconnect synthesis was not possible prior to this work, due to the inefficiency of the previous buffering techniques.

After the publication of this work, various extensions to the slew buffering-driven technique were developed by other experts in the field. In summer 2010, I was invited by the group again to take a visiting professorship working on physical design, which resulted in a US patent being granted. 

## PRODUCT NEWS

### I/O EXPANSION FOR RUGGED APPLICATIONS

The **SBC35-CC405** series of multi-core embedded PCs includes on-board USB, gigabit Ethernet (GbE), and serial ports. These computers are designed for rugged embedded applications that require extended temperature operation and long-term availability.

The SBC35-CC405 series features the latest-generation Intel Atom E3800 family of processors in an industry-standard 3.5" single-board computer (SBC) format COM Express carrier. A Type 6 COM Express module that supports a quad-, dual-, or single-core processor is used to integrate the computer.

For networking and communications, the SBC35-CC405 multi-core embedded PCs include two Intel I210 GbE controllers with IEEE 1588 timestamping and 10-/100-/1,000-Mbps multispeed operation. Four Type-A connectors support three USB 2.0 channels and one high-speed USB 3.0 channel. Two serial ports support RS-232/-422/-485 interface levels with clock options up to 20 Mbps in the RS-422/-485 mode and up to 1 Mbps in the RS-232 mode.

The SBC35-CC405 series also includes two Mini PCI Express (PCIe) connectors and one IO60 connector for additional I/O expansion. Both Mini PCIe connectors support half-length and full-length cards with screw-down mounting for improved shock and vibration durability. One Mini PCIe connector also supports bootable mSATA solid-state disks while the other connector includes USB.

The IO60 connector provides access to the I<sup>2</sup>C, SPI, PWM, and UART signals, enabling a simple interface to sensors,



data acquisition, and other low-speed I/O devices.

The SBC35-CC405 runs over a 10-to-50-VDC input power range and operates at a -40°C-to-85°C temperature range. Enclosures, power supplies, and configuration services are also available.

Linux, Windows, and other x86 OSES can be booted from the CFast, mSATA, SATA, or USB interfaces, providing flexible data storage options. WinSystems provides drivers for Linux and Windows 7/8 as well as preconfigured embedded OSES.

The single-core SBC35-CC405 costs **\$499**.

**WinSystems, Inc.**  
[www.winsystems.com](http://www.winsystems.com)

### LONG-RANGE RFID IC

The **EM4126** EPC radio-frequency identification (RFID) IC is designed to provide RFID tagging on small and/or high-value products (e.g., jewelry and watches).

The IC's high sensitivity enables long read ranges. EM4126-based tags can achieve -21-dBm read sensitivities. The ICs are designed for supply chain management, tracking and tracing, container identification, and access and asset control applications.

The EM4126 RFID IC includes 224 bits of nonvolatile memory, which supports International Organization for Standardization (ISO) or Electronic Product Code (EPC) data structures. The memory also enables SGTIN-198 encoding, which uses alphanumeric serialization represented as a string of up to 20 7-bit characters.

The RFID IC's additional features include ISO 18000-63 and EPC Class-1 Generation-2 compliance, 32-bit short-tag

identification, 40-to-160-Kbps forward- and return link-data rates, and a -40°C-to-85°C extended temperature range.

Contact EM Microelectronic for pricing.

**EM Microelectronic**  
[www.emmicroelectronic.com](http://www.emmicroelectronic.com)



## PRODUCT NEWS


### HIGH-ELECTRON-MOBILITY TRANSISTORS

The **TPH3002LD** and the **TPH3002LS** are 600-V gallium nitride (GaN)-based, low-profile power quad flat no-lead (PQFN) high-electron-mobility transistors (HEMTs). The HEMTs utilize Transphorm's patented, high-performance EZ-GaN™ technology, which combines low switching and conduction losses, reducing the overall system energy dissipation up to 50%.

The TPH3002PD and TPH3002PS HEMTs are designed for use in smaller, lower-power applications (e.g., adapters and all-in-one computer power supplies). The HEMTs feature a Kelvin connection to isolate the gate circuit from the high-current output circuit to further reduce electromagnetic interference (EMI) and high-frequency switching capabilities.

Evaluation boards are available with the devices. Contact Transphorm for pricing.

**Transphorm, Inc.**  
[www.transphormusa.com](http://www.transphormusa.com)




### 8-BIT MICROCONTROLLER IP CORE

The **DF6808** IP core is binary-compatible with the industry-standard Motorola 68HC08 8-bit microcontroller. The IP core uses sophisticated on-chip peripheral capabilities to perform 45 to 100 MIPS. FAST architecture implemented in DF6808 enables the 68HC08 microcontroller to run at least three times faster than the original solution.

The DF6808's 16-bit, free-running timer system has two input-capture lines and two output-compare lines. The IP core is equipped with proprietary safety functions including self-monitoring circuitry, which helps protect against system errors; the computer operating properly (COP) watchdog system, which protects against software failures; and an illegal opcode detection circuit, which provides a non-maskable interrupt if an illegal opcode occurs.

For power conservation, the IP core includes two software-controlled power-saving modes (Wait and Stop). These modes make the DF6808 IP core well suited for automotive and battery-driven applications.



The DF6808 includes the DoCDTM real-time hardware debugger, which provides built-in support for Digital Core Design's hardware debug system and the debugging capability of an entire system-on-a-chip (SoC). The DoCDTM enables nonintrusive debugging of running applications. It can halt, run, step into, or skip an instruction and read/write any microcontroller contents, including all registers, user-defined peripherals, data, and program memories.

Contact Digital Core Design for pricing.

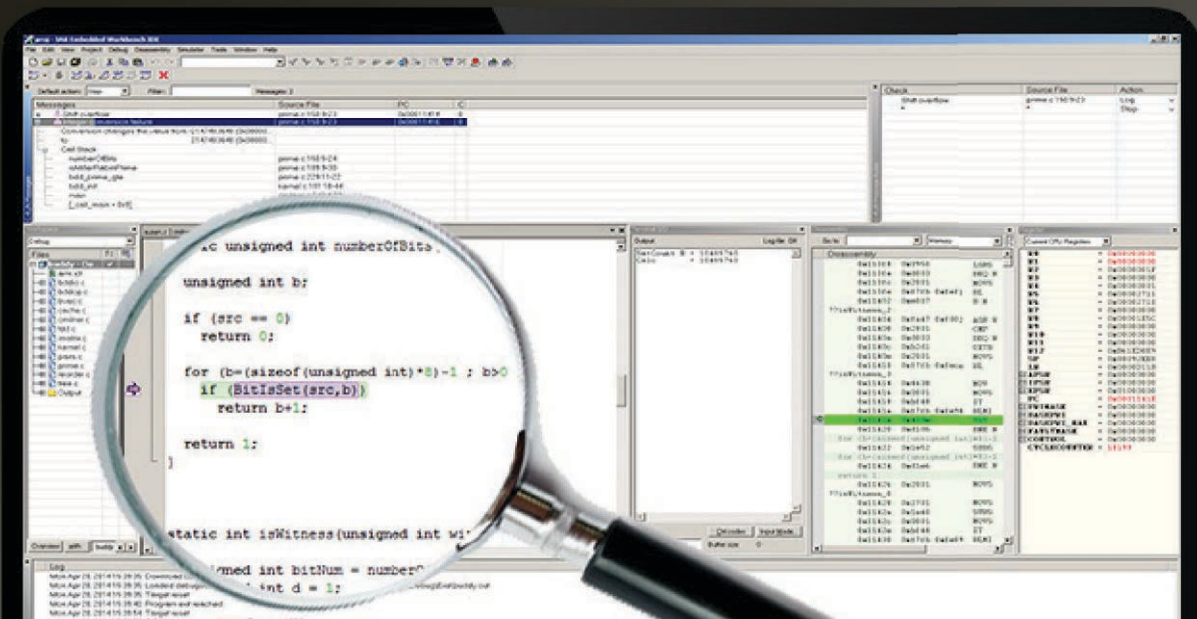
**Digital Core Design**  
[www.dcd.pl](http://www.dcd.pl)



# EXPLORE

# C-RUN FOR ARM

## Runtime Analysis Simplified



C-RUN is a high-performance runtime analysis add-on product, fully integrated with world-leading C/C++ compiler and debugger tool suite IAR Embedded Workbench.

C-RUN performs runtime analysis by monitoring application execution directly within the development environment.

The tight integration with IAR Embedded Workbench improves development workflow and provides each developer with access to runtime analysis that is easy-to-use.

## PRODUCT NEWS

### LOW-POWER MICROMODULE

The **ECM-DX2** is a highly integrated, low-power consumption micromodule. Its fanless operation and extended temperature are supported by the DMP Electronics Vortex86DX2 system-on-a-chip (SoC) CPU. The micromodule is targeted for industrial automation, transportation/vehicle construction, and aviation applications.

The ECM-DX2 withstands industrial operation environments for -40-to-75°C temperatures and supports 12-to-26-V voltage input. Multiple OSes, including Windows 2000/XP and Linux, can be used in a variety of embedded designs.


The micromodule includes on-board double data rate (DDR2) memory that supports up to 32-bit, 1-GB, and single-channel 24-bit low-voltage differential signaling (LVDS) as well as video graphics array (VGA) plus LVDS or VGA plus TTL multi-display configurations.

The I/O deployment includes one SATA II interface, four COM ports, two USB 2.0 ports, 8-bit general-purpose input/output (GPIO), two Ethernet ports, and one PS/2 connector for a keyboard and a mouse. The ECM-DX2 also provides a PC/104 expansion slot and

one Mini PCI Express (PCIe) card slot.

Contact Avale Technology for pricing.

**Avale Technology, Inc.**  
[www.avale.com.tw](http://www.avale.com.tw)




### 3-D PCB SOFTWARE

The **brd-to-3D** is a virtual 3-D EAGLE viewer designed to work with Beta LAYOUT's PCB-POOL PCB prototyping module. The 3-D tool complements CadSoft's EAGLE software that currently does not have a direct STEP export.

When you upload the EAGLE layout file, you receive a "free" comprehensive 3-D package that includes photorealistic images of your PCB and surface-mount device (SMD) stencil, a STEP file of the virtually assembled PCB, and a PDF-formatted 3-D view.

A free laser-sintered 3-D model of the assembled PCB is available with the purchase of a PCB-POOL prototype order. Contact Beta LAYOUT for pricing.

**Beta LAYOUT, Ltd.**  
[www.beta-layout.com](http://www.beta-layout.com)



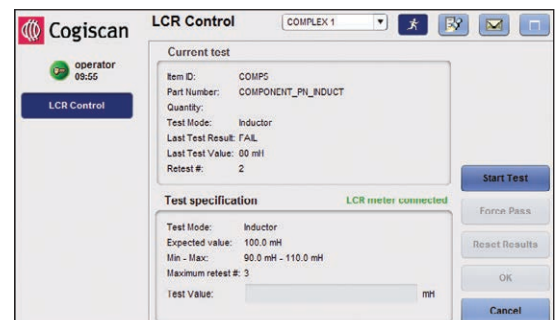
### CLOSED-LOOP MODULE

The **LCR Control Module** is a fully integrated and closed-loop system designed to eliminate the risk of placing incorrect passive components on PCBs. Combined with the Offline Job Setup and/or Line Setup Control module, the LCR Control Module's software is integrated with an electrical LCR meter.

The system verifies that the measured values (e.g., inductance, capacitance, and resistance) are within the tolerances specified for the component. This electrical verification can be done at any time prior to mounting the reel on the placement machine.

Contact Cogiscan for pricing.

**Cogiscan, Inc.**  
[www.cogiscan.com](http://www.cogiscan.com)



# Get ready for



# Buy now and upgrade for free!

Valid for new license purchases only.

Terms & Conditions: [www.cadsoftusa.com/conditions-freeupgrade](http://www.cadsoftusa.com/conditions-freeupgrade)

[www.cadsoftusa.com](http://www.cadsoftusa.com)




## PRODUCT NEWS

### EMC MEASUREMENT TECHNOLOGY

SX near-field probes enable electromagnetic compatibility (EMC) analysis of interferences emitted by electronic boards, components, and IC pins with high internal frequencies. The **SX-R3-1** magnetic H-field probe is designed to detect high-frequency magnetic fields with a high geometrical resolution.

The field orientation and distribution can be identified by moving the SX-E03 E-field probe around conductor runs, bypass capacitors, EMC components, and within IC pin and supply system areas. The probe detects bus structures and larger components.

The probes have a 1-to-10-GHz frequency range. Their high resolution (the SX R3-1 achieves 1 mm and the SX E03 covers up to 4 mm × 4 mm) enables them to pinpoint RF sources on densely packed boards or on IC pins. The magnetic-field probe heads are electrically shielded. The probes are connected to a spectrum analyzer input via a shielded cable and SMA connectors during measurement. High clock rates of 2 GHz, for example, may result in fifth-order harmonics of up to 10 GHz. These harmonics are coupled out by RF sources on the board (e.g., conductor-run segments, ICs, and other components). They may stimulate other structural parts of the board to oscillate and emit interferences.



Contact Langer EMV-Technik for pricing.

**Langer EMV-Technik**  
[www.langer-emv.com](http://www.langer-emv.com)

### INTEGRATED WI-FI SYSTEM IN PACKAGE MODULE

The **EC19W01** is a small, smart, highly integrated 802.11b/g/n Wi-Fi system in package (SiP) module. The module is well suited for home automation and smart appliances; Wi-Fi audio speakers and headphones; wireless sensors and sensor networks; wireless monitoring (audio and video); smart appliances; health care and fitness devices; wearable devices; security, authentication, and admittance control; lighting; building/energy/industrial management/control; cloud-connected devices; remote control, data acquisition, and monitoring; and machine-to-machine (M2M) and Internet of Things (IoT) design.

The EC19W01's features include an integrated 32-bit processor that supports application customization, on-board flash memory and an antenna, low power consumption, support for Serial-to-Wi-Fi and SPI-to-Wi-Fi, wireless transmit/receive rates of up to 20 Mbps, and a small 14-mm × 16-mm × 2.8-mm footprint.

Contact Econais for pricing.

**Econais, Inc.**  
[www.econais.com](http://www.econais.com)



**CLIENT PROFILE**

# TechTools

[www.tech-tools.com](http://www.tech-tools.com)

P.O. Box 2408 Rowlett, TX 75030-2408

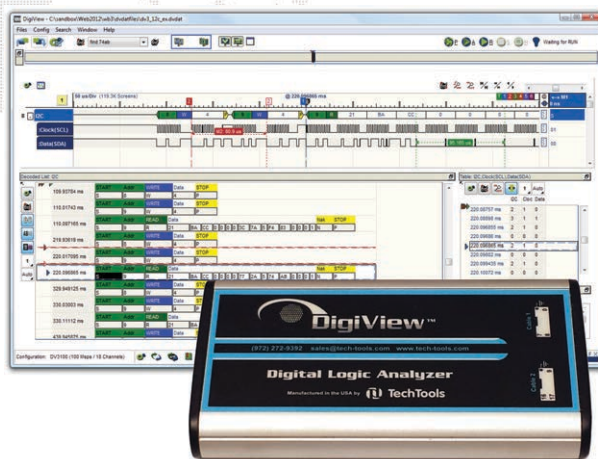
**CONTACT:** sales@tech-tools.com

**FEATURED PRODUCT:** The DigiView series of logic analyzers are real, full-featured affordable instruments. DigiView hardware performs all real-time functions including sampling, triggering, compression, and storage. The PC's computing power and resources provide an enhanced data display, analysis, and sharing.

This division of labor results in the same sample rate, capturing, triggering, and storage performance of expensive stand-alone instruments. Leveraging the PC also enhances data handling and saves costs.

All DigiView models incorporate advanced, real-time, adaptive compression for full resolution captures and long capture times. Prices start at \$349. Models are available with 100 to 400 MSPS and nine to 36 channels. For full details, visit [www.tech-tools.com/real3.htm](http://www.tech-tools.com/real3.htm).

**EXCLUSIVE OFFER:** During the month of July, *Circuit Cellar* members will receive 10% off the purchase of any DigiView



model. To qualify, first place an order for a DigiView online at [www.tech-tools.com/real3.htm](http://www.tech-tools.com/real3.htm) by July 31, 2014. Then e-mail a copy of your order acknowledgment to sales@tech-tools.com and mention PROMO CODE: CCIREAL3. A 10% discount will be applied to your order total.


*Circuit Cellar prides itself on presenting readers with information about innovative companies, organizations, products, and services relating to embedded technologies. This space is where Circuit Cellar enables clients to present readers useful information, special deals, and more.*

INDUSTRY & ENTERPRISE

## Industrial Temperature SBC

### iPac-9x25

- \*Atmel 400 Mhz Processor
- \*128MB DDR2 RAM, 4GB eMMC
- \*16MB Serial Data Flash, Micro SD
- \*36 General Purpose Digital I/O lines
- \*8 High Drive Digital Outputs
- \*1x USB 2.0 (High-Speed) Host
- \*1x USB (Full-Speed) Host
- \*1x USB 2.0 (High-Speed) Device Port
- \*3x RS232, 1x RS232/422/485
- \*2x 10/100 Ethernet, 1x CAN Bus
- \*1x I2C Port & 1x SPI Port
- \*Up to 7 channels of 10 bit A/D
- \*Up to 4 16-bit PWMs
- \*Industrial Temp. -40C to +85C



Designed and manufactured in the USA, the iPac-9X25 is a web enabled microcontroller with the ability to run an embedded server and to display the current monitored or logged data. The web connection is available via two 10/100 Base T Ethernet ports or 802.11 wireless wifi networking when using the proper Linux modules and adapters. This Microcontroller has all connectors brought out as headers on a board and has the same footprint of a standard PC/104 module at 3.77" x 3.54". The iPAC-9X25 is perfectly suited for Industrial Temperature Embedded Data Acquisition and Control applications. Pricing for Qty 1 is \$198

[www.emacinc.com/sales/ipac14](http://www.emacinc.com/sales/ipac14)

Since 1985  
OVER  
29  
YEARS OF  
SINGLE BOARD  
SOLUTIONS

# EMAC, inc.

EQUIPMENT MONITOR AND CONTROL

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

## SCOPE DEALS

### PASSPORT-SIZE PC SCOPES

Great scopes for field use with laptops. Up to 200MHz bandwidth with 1GSa/s, high speed data streaming to 1MSa/s, built-in 1GSa/s AWG/wfm gen. - PS2200A **\$239+**



### 30MHz SCOPE

Remarkable 30MHz, 2-ch 250MS/s sample rate oscilloscope. 8-in color TFT-LCD and AutoScale function. Includes FREE carry case + 3 yr warranty! - SDS5032E **\$289**



### 60MHz SCOPE

Best selling 60MHz 2-ch scope with 500MSa/s rate + huge 10MSa memory! 8" color TFT-LCD. Includes FREE carry case! - SDS6062 **\$349**



### 100MHz SCOPE

High-end 100MHz 2-ch 1GSa/s benchscope with 1MSa memory and USB port + FREE scope carry case. New super low price! - DS1102E **\$399**



### 70-300MHz SCOPES

Fast, versatile 2-ch 2GSa/s scopes with 8" WVGA LCD, integrated generator, 14Mpt memory, very low noise floor. - DS2000A series **\$839+**



**INCREDIBLE LOW PRICES, FREE TECHNICAL SUPPORT  
GREAT CUSTOMER SERVICE SAELIG.COM**

# Personal Solar Power Setup (Part 1)

## Solar Panels, Batteries, and Storage



This article series describes a solar power system designed to provide reliable backup power and reduce utility power usage. Part 1 focuses on the design plans, including the panels and the power source.

By Steve Hendrix (US)

Everywhere you turn these days, someone wants to sell you a “green-energy solution.” Where I live, in Ohio, utility power is fairly reliable and reasonably inexpensive. When we have an outage, it’s usually either a “blink” or a really extended outage stretching to hours or days. I have a standby generator, but it’s noisy and smelly and it burns gas. Solar sounds much better and can ease reliance on utility power. My two main reasons for adding solar power are to have reliable backup power during an outage and to “go green” to reduce utility power usage.

I have been a die-hard follower of *Circuit Cellar* since the *Byte* magazine days, so of course I couldn’t go out and buy a prepackaged solution, even if price were no object. However, I learned that the real-world issues of actually implementing a complete solution go far beyond the solar panels’ appeal. The mundane issues of physical mounting, wiring, system protection, energy storage, and conversion to usable power took far more effort than purchasing the panels.

There were many tradeoffs to consider as I designed my system. The first part of this article series will explain my reasoning and the solutions I used for the “grunt-work” parts of the project. In Part 2, I will delve into the electronics and the issues involved with mixing utility power with solar.

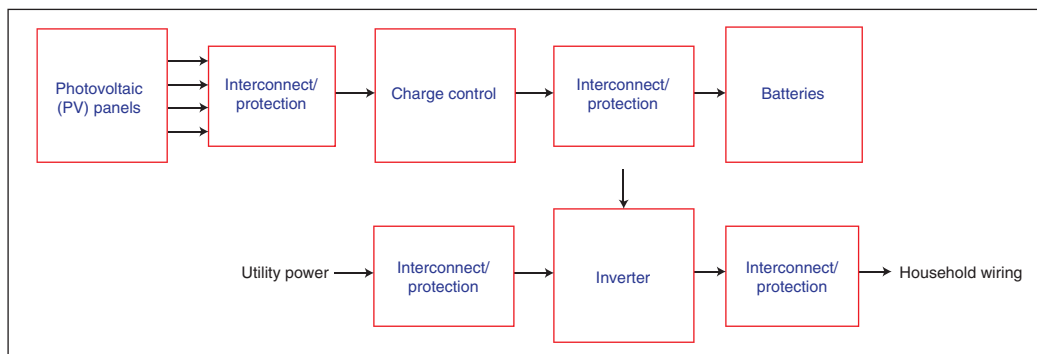
You will hear about safety issues throughout the article series. I will cover the basics, but check your local electrical codes, as they differ. I may be rather paranoid about safety. I get that way when my family has to sleep with the system running. I’ve worked on live electronics a lot, and sometimes I have been bitten hard by a wire that I “knew” was “off.”

### SOLAR PANELS AND MOUNTING ISSUES

When you think of solar power, you likely picture a big array of solar cells facing the sun, so I’ll start with that part of the system. There are several interrelated issues here. **Photo 1** shows a portion of my system mounted on a south-facing roof with a standard 5:12 pitch. **Figure 1** shows the system’s overall block diagram, for reference in the rest of this article series.

#### PHOTO 1

Solar panels mounted on a south-facing roof are shown. Note the 4” gap under the bottom of each. This enables rainwater, leaf litter, and balls to pass under rather than accumulating. The upper brackets on each panel are sized to hold the panels at 50° to horizontal and the rows are spaced to avoid shading one another even during the winter solstice.



**FIGURE 1**

The personal solar power system includes batteries, photovoltaic (PV) panels, interconnect protection, wiring, and power. Interconnections and protection are a very important part of the system.

In Part 1, I will focus mainly on the top half of **Figure 1**. Part 2 will explore the details of the bottom half and go into more depth about charge control.

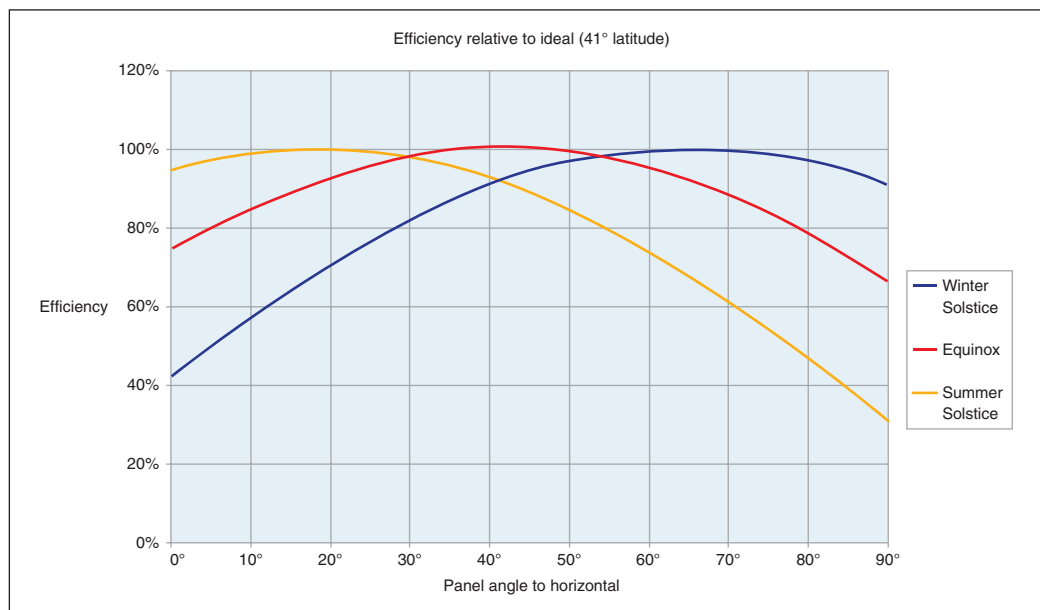
First, I needed to estimate the total number and area of panels needed for the load I wanted to power. An effective 6 h of direct sun per day accounts for the fall off toward sunrise and sunset. In Ohio winters, the average effective sun will be a lot less! Thus a 200-W panel can provide around 1,200 watt-hours (WH) of energy per day under good conditions. If there is enough storage to carry through the night, this runs a continuous 50-W load.

When I began looking into solar, panels cost about \$6 per watt. As prices dropped, I placed a bulk order with several other people to get the price down to \$1.80 per watt. I wanted to optimize the amount of solar power collected for a given investment. In an ideal world, the panels would squarely face the sun. The amount of power collected falls off with the cosine of the angle between the incident ray and the panel's plane. Daily movement requires motors and moving parts

and is prone to obstacles, wind interference, and so forth. Thus, most installations I've seen opt for no moving parts. The actual loss is less than you may expect, because the light missed near sunrise and sunset comes through a less-transparent atmosphere. However, the north-south tilt deserves some thought. At the equinox, the panels' optimum slope matches the latitude. This angle would seem to optimize the system throughout the year.

I chose to set the panels steeper for two reasons. First, if the load is roughly the same year-round, I care a lot more about optimization in the winter, when there are many cloudy days and much less daylight than in the summer, with its multitude of sunny days and more overall daylight. Second, I live in snow country. A few inches of snow is common, and I've had a couple feet of snow on the roof. Thus I wanted a higher angle to shed snow more quickly.

**Figure 2** shows that for my 41° latitude the system loses a little efficiency in the summer, but it gains winter efficiency and sheds snow much better with my chosen 50°




**FIGURE 2**

This graph shows solar panel efficiency changes with the angle of the sun. The year-round optimum angle matches the local latitude. However, for improved efficiency in winter and better snow shedding, I chose a higher 50° angle. This angle yields 97% at the winter solstice, 99% at both equinoxes, and 85% at the summer solstice.

# Elektor.STORE

The world of electronics  
at your fingertips!



**NEW!**


This book and more  
are available at  
[www.elektor.com/store](http://www.elektor.com/store)

The Ultimate Guide!

## The LTSpice IV Simulator

LTSpice IV is one of the most extensive SPICE simulators currently available, with over 1.5 million users worldwide. It is completely free, allowing you to work with an unlimited number of nodes, components and simulations. So much so, that many of the features on LTSpice IV continue to go unexplored and unexploited and, unfortunately, literature and technical documentation are hard to come by. Recently made available in the Elektor Store, The LTSpice IV Simulator book is the ultimate guide to LTSpice IV. In a sturdy, hard-cover format, The LTSpice IV Simulator describes the operation of the program, all available commands, the various editors, dealing with SPICE models, the use of non-linear components and more. Perhaps the most robust piece of literature on the topic, The LTSpice IV Simulator is more than just a manual. It also offers a variety of tips, methods and examples, all carefully illustrated using almost 500 drawings, diagrams and screenshots on high-quality paper. The book is designed so that it is suitable for both beginner and veteran SPICE users.

744 pages • ISBN 978-3-89929-258-9 • \$67



**Bestseller!**

Learning to Fly With Eagle

## Eagle V6 Getting Started Guide

The book is intended for anyone who wants an introduction to the capabilities of the CadSoft's Eagle PCB design software package. The reader may be a novice at PCB design or a professional wanting to learn about Eagle, with the intention of migrating from another CAD package. After reading this book while practicing some of the examples, and completing the projects, you should feel confident about taking on more challenging endeavors. This book is supplied with a free copy of Eagle on CD-ROM for MS Windows, Linux and Mac.

208 pages • ISBN 978-1-907920-20-2 • \$47



The RPi in Control Applications

## Raspberry Pi Hardware Projects

This book is about the Raspberry Pi computer and its use in control applications. Dogan Ibrahim explains in simple terms, with examples, how to configure the RPi, how to install and use the Linux operating system, how to write programs using the Python programming language and how to develop hardware based projects.

290 pages • ISBN 978-1-907920-29-5 • \$54






### MIFARE and Contactless Cards in Application **RFID**

MIFARE is the most widely used RFID technology, and this book provides a practical and comprehensive introduction to it. Among other things, the initial chapters cover physical fundamentals, relevant standards, RFID antenna design, security considerations and cryptography. The complete design of a reader's hardware and software is described in detail. The reader's firmware and the associated PC software support programming using any .NET language.

484 pages • ISBN 978-1-907920-14-1 • \$68



### A Highly Practical Guide **Process Measurements with C# Applications**

This book introduces PC based measurement systems and software tools for those needing to understand the underlying principles or apply such techniques. C# was chosen to give readers immediate, practical experience with a prominent programming language. C-Sharp has a wide support base and is a popular choice for engineering solutions. The basics of measurement and data capture systems are presented, and examples of software post-processing.

144 pages • ISBN 978-1-907920-24-0 • \$39



### 110 Elektor Editions, Over 2500 Articles **DVD Elektor 2000 through 2009**

This DVD-ROM contains all circuits and projects published in Elektor magazine's year volumes 2000 through 2009. The 2500+ articles are ordered chronologically by release date (month/year), and arranged in alphabetical order. A global index allows you to search specific content across the whole DVD. Every article is printable using a simple print function. This DVD is packed with ideas, circuits and projects that are ideal for any electronics enthusiast, student or professional, regardless of whether they are at home or elsewhere.


ISBN 978-1-907920-28-8 • \$121

**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/store](http://www.elektor.com/store)



Elektor US  
111 Founders Plaza, Suite 300  
East Hartford, CT 06108  
USA  
Phone: 860.289.0800  
Fax: 860.461.0450  
E-mail: [order@elektor.com](mailto:order@elektor.com)




### **Android Breakout Board**

The FTDI FT311D is a flexible bridge that can interface your circuit to an Android smartphone or tablet. This Elektor Android Breakout Board offers options for seven digital outputs, four PWM outputs, asynchronous serial and I2C and SPI interfaces. The board is compatible with Android 3.1 (Honeycomb) or higher (Android Open Accessory Mode should be supported).

*Ready-built module*

Art.# 130516-91 • \$41




### **E-Lock**

The E-Lock chip allows you to connect your control system to the 'Network of Networks' and monitor and control it from anywhere on or around the globe using your computer, tablet or smartphone. This completely assembled module is able to establish a secure connection over the Internet (TCP/IP) using the Transport Layer Security (TLS) encryption protocol. The chip also offers a 7-line GPIO (general purpose input/output), four 16-bit ADC channels, one 12-bit DAC and one I2C bus that allow control and communication with other peripheral devices.

*Ready-built module*

Art.#130280-91 • \$150



**FIGURE 3**

These are the mounting details and bracket dimensions. Lower brackets hold the panels up off the roof to avoid trapping debris underneath. Upper brackets hold the panels at the calculated angle to horizontal to optimize year-round efficiency. Row spacing is calculated to prevent a lower row from shading an upper row at the winter solstice, when the sun is lowest in the sky.

angle to horizontal. (See the spreadsheet in the references and adjust for your latitude.) After 2013's grim Ohio winter, I'd go up to 60° if I were doing it over.

Next, I needed to determine how to mount the panels. A south-facing roof at the selected angle would be ideal, although at 50° it would be a little awkward to work on. If you are roof mounting, there needs to be room for air circulation underneath. Don't lay panels



**PHOTO 2**

The bottom edge of the panel runs straight down the roof's slope so drainage is not an issue. But the upper brackets must be longer to maintain the desired angle to the sun. There is a quarter-twist in the brackets so the feet face uphill with the end flush against the edge of a shingle. This way water and debris cannot collect under the curve of the bracket against the roof.

flat on shingles, as the heat impairs their efficiency.

My solution was to buy 0.25" x 1" aluminum bar stock and cut it to the required length. The bar bent easily and it was straightforward to drill screw holes. On my south-facing roof, I made the brackets to hold the panel's bottom edge about 4" off the roof, with the upper brackets long enough to give me the chosen angle. The gap at the bottom enables snow to slide off, but more importantly lets leaves, balls, and other debris slide right under rather than accumulating. **Figure 3** shows a side view of my mounting calculations (seen from the west in the northern hemisphere or from the east in the southern hemisphere).

Even if your sloping roof does not face directly south, you needn't abandon roof mounting. I needed more space than I had on a south face, so I installed some panels on an east-facing roof. On my relatively flat 5:12 roof slope, the only noticeable difference is that optimum sunlight happens earlier in the day. However, I did have to use longer upper brackets to achieve the same angle to horizontal. In return, I used very small lower brackets since the lower edge now runs directly down the slope, thus nothing gets caught under it. That enabled me to feel comfortable with 0.125" bar stock there. **Photo 2** shows details of the mounting on the east-facing roof.

You must also consider wind. Typical panels are big! Don't carry them around on the roof in a high wind. More importantly, you will need to ensure the brackets are well-secured to the underlying structure. Shortly after mine were up without everything fully secured, we had hurricane remnants that resulted in 60-knot winds out of the north, which was absolutely the worst possible setup for my panels. Fortunately, I was able to get back up on the roof and secure most of the brackets before the storm hit. Everything came through fine, but I had a nervous night. To better secure the brackets, I used some of the leftover aluminum pieces to lay right across the bracket feet. This enabled me to spread out any pulling forces across many more screws and more roof area.

If I were doing it over, I would at least consider a ground mounting. I could erect a lean-to structure of wood or steel beams. The panels would bolt to the structure and provide the south half of the roofing. This would enable easier access to the wiring and less up-high work, as well as potentially providing a storage shed. However, this solution has the drawbacks of being more readily shaded by trees, damaged by people or animals, and, of course, weeds growing up underneath.

The biggest safety issues are simple

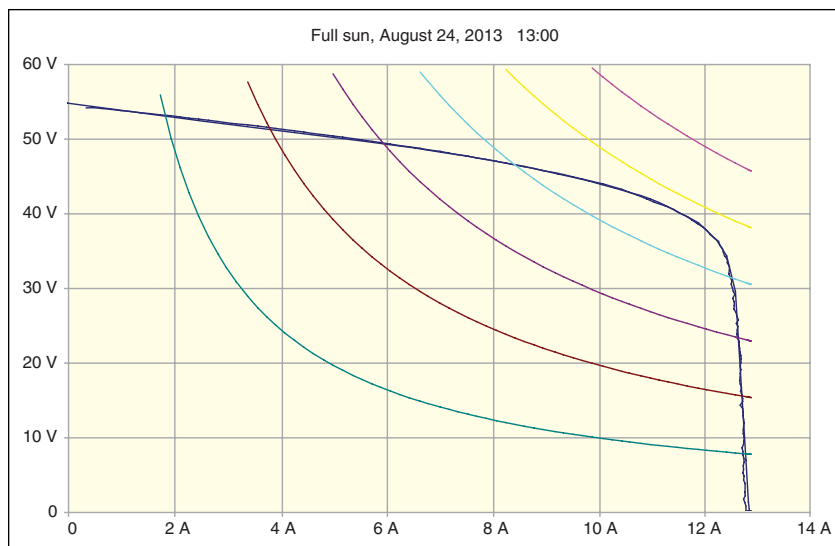
things (e.g., falling off a roof, having a panel blown out of your hands, etc.). The electrical safety issue with a single panel or even a series string is actually fairly minimal, even in the sunshine. The panels I used (they seem to be among the most common) are rated for about 200-W maximum power, about 21-V maximum open-circuit voltage, and about 12-A maximum short-circuit current. I used a KEPCO electronic load to make measurements and generate the curve shown in **Figure 4**.

There are three panels in series showing how they really behave. Realistically, there are only two conditions you're likely to apply by accident. First, an open circuit (or nearly so) by skin touching the two terminals: 21 V won't even give you a good tickle. The second likely possibility is a wrench or a screwdriver across the terminals: 12 A at 0 V isn't going to do much of anything exciting either. So unless you string a lot of panels in series, the electrical risks of working on them are fairly minor. But working at night or twilight pretty much negates the electrical risk, as long as you have enough lighting to see what you're doing.

## BATTERIES

There are several new battery technologies, but plain-old lead-acid batteries provide by far the best bang for the buck (deep-cycle type preferred). The most readily available and economical large-capacity batteries are 6-V golf-cart batteries. You need to connect these batteries in series strings to achieve the desired system voltage, and you may or may not need several strings in parallel to achieve the required current and/or total energy storage requirements. Don't scrimp on this wire. You should use automotive starter-style wire (4 AWG or heavier) and terminals as you may have hundreds of amps going in or out.

Before going further, I must provide a safety warning about batteries. Unlike the solar panels, which are fairly difficult to provoke electrically into an unsafe situation, big lead-acid batteries can be quite dangerous. There are the usual cautions about chemical burns as well as the simple physical strains of their weight (75 lb apiece for my big golf-cart batteries). The batteries deliver a lot of current at full voltage. A small car battery can deliver up to about 20x its rated amp-hour (AH) capacity. A 50-AH battery will only have a minor voltage drop if asked to deliver 1,000 A. These deep-cycle batteries aren't quite as efficient with sudden heavy loads, but nonetheless can deliver quite substantial currents. Thus, before I connect a wire to a battery, I ensure the other end is already solidly connected to its final resting place and



**FIGURE 4**

This is a voltage-current curve for a series string of three 200-W panels. The open circuit (0-A) voltage is about 55 V. The short circuit (0-V) current is about 13 A. The maximum power is near 11.5 A and 41 V. The concave curve family marks lines of constant power at 100–600 W.

at least one wire has some kind of protective device. And of course I'm very conscious of exactly where any metal tools are around those terminals. I don't like to make little crescent-wrench puddles!

The charging algorithm for lead-acid batteries is fairly simple. Note that any such battery is made up of several 2-V cells in series. When charging, you can pretty much give it all the current you have available until the voltage reaches 2.41 V per cell. This is termed "bulk charge." Then the charger backs off the current enough to hold this voltage for 1 to 2 h, depending on the battery, which is termed "soak charge." Finally there is the "float charge" state, where the current is shut off, turning it back on enough to maintain 2.35 V per cell.

Any current being drawn off by a load is irrelevant to the controller, as long as it can maintain the desired voltages. During discharge, you want the voltage to go no lower than 1.75 V per cell.

Because of inverter availability as well as a lot of common automotive parts, a 12-V system was my first choice. As part of a group buy with several friends, I purchased a pair of 60-A charge controllers and set up my system as two almost completely independent systems with five panels in parallel on each. Therefore, I had available a nominal 18 V at 55 A on each of the two systems.

For the batteries, I set up the two systems each with three parallel pairs of two series-connected 6-V, 220-AH batteries. Thus each system had nominally 12 V at 660 AH available battery storage, for a total of 1,320 AH. The inverter I used is one of the bigger consumer-grade inverters, which provide 120 VAC at 4,500 W continuous and 9,000 W peak.

When I expanded the system and needed to install panels further from my battery

bank, I computed the cost of the wire and the energy lost in the wires and concluded that I needed to go to a higher system voltage. I used 48-V inverter with a standard utility-style split 240-/120-VAC output and some special features, which I will discuss in Part 2 of this article series.

Then I added my 240-VAC deep-well pump in addition to other essential loads (e.g., my refrigerator, deep freeze, and furnace blower and controls). I wouldn't want to try running the air conditioning compressor too. Last summer I accidentally left the furnace blower on solar and discovered that it is actually about half the electrical load of my air conditioning. And now and then I even have the well pump and the air conditioning running at the same time. This provides a pretty impressive current spike from the batteries when they kick on!

For the 48-V system, I reasoned that a series string of three panels should have a nominal maximum power point at 54 V at 11 A and the rated 21-V open-circuit voltage would yield a maximum of 63 V under no-load conditions. With a nominal 48-V battery, you want bulk charging to end at 58 V. Float charge should happen at 54 V. Thus, with even a minimal load, the system voltage shouldn't be able to go much above the ideal 58-V charging voltage and when pulled down to 48 V with the batteries very discharged, there should be plenty of current. This is the ultimate "poor man's charge controller!"

Unfortunately, things are never that simple. It worked fine in the late winter and early spring, but once we had a run of several nice clear days of summer sun, and the batteries reached full charge, the voltage rose high enough that I manually turned off some panels in the afternoon on bright sunny days.

## INTERCONNECTIONS

Based on decades of electronics experience, I know that connections are the failure points in any electrical system, so I planned for connections up front. The solar panels come with pre-installed connectors. Fortunately, the industry has become somewhat standardized, so much that electronic component distributor Digi-Key even lists a special category under connectors for "photovoltaic (PV) connectors." Two big drawbacks of these connectors are that they need a special crimp tool to attach to a wire and they require a special tool to unlock them once they're plugged together. I've never liked crimp connections, especially for high current, so I used a standard crimper on them and then soldered it anyway. I prefer molten metal for connections that have to stay reliable.

As for the locking mechanism, you can use an ordinary screwdriver to unlock it, if needed. The connectors will make contact for testing when partially plugged together but not locked. Two important points: Check and recheck polarity so you get the male and female connectors on the correct wires and pay attention to the order things must slide up the wire *before* you solder the contact.

I use commodity wiring for long wire runs. Romex wire comes in a variety of appropriate sizes, conductor counts, and insulation types. Most Romex wire these days includes a bare ground wire. I used that to ground the metal frames of the panels to a common ground point. Now I will discuss wire losses.

Wire gauge works much like a decibel scale for resistance per foot. 0 AWG is 100  $\mu\Omega$  per foot. Every 10 steps of gauge is a factor of 10, 20 steps is a factor of 100, three steps is a doubling, and six steps is 4x. Thus for 14 AWG, I would start at 20 AWG at 10 m $\Omega$  per foot, and go down by 6 dB (a factor of 4) to find 2.5 m $\Omega$  per foot for 14 AWG. For 10 AWG, it's a nice simple 1 m $\Omega$  per foot.

Thus if I want to use 14 AWG wire favored for inexpensive household wiring, I can expect 2.5 m $\Omega$  per foot. So, for an 80' run (160' round trip), you would have 0.4  $\Omega$  in the wiring. If you are to run 11 A through that, you lose 4.4 V. That may be an acceptable percentage loss at 120 V, but at 12 V, you lose more than a third of the power by the other end of the wire. This is why power companies prefer to transmit power over long distances at a high voltage and a low current.

This wire resistance issue, along with cost, caused me to run most of my panels as series strings of three panels and to use 10-2 Romex for those runs. The furthest panels needed an 80' run, but since I used 10 AWG wire, the total round-trip resistance was only 0.16  $\Omega$ , which caused a 1.76-V drop at 11 A. A voltage of 1.76 out of 48 V represents a much more acceptable 3% loss. The closer panels have shorter wire runs and correspondingly lower losses.

I could, of course, have done like some of my friends, and placed junction boxes on the roof. Such junction boxes reduce the total number of wires but require much heavier wires and correspondingly heavier connections, switches, and so forth. The bottom line is that I can disconnect a faulty panel or series string, snug and warm in my garage with tools handy and good lighting, rather than up on the roof in the snow or rain.

As part of a professional project a while back, I needed a test stand switch that could handle about 100 A of 12 VDC. They were quite expensive. I also needed some kind of enclosure, wire terminations, and protection.

I was doing a minor upgrade to my home wiring when it hit me: A plain old household circuit breaker and panel would do the entire job with no customization and all the needed parts are readily available and inexpensive, due to their commodity nature. I have since standardized with one brand and I have dozens of those boxes. I reasoned that the circuit breaker can't see the neutral wire in a standard installation, so it doesn't care about the voltage for which it's being used (as long as it can handle the voltage if it opens). I later did some tests and found data confirming that these particular breakers don't care whether they are passing AC or DC. My tests showed a thermal trip that takes about 30 s at a 2x rating and a magnetic (near-instant) trip at a 10x rating. You will need to check the data for a particular circuit breaker to ensure its ratings are adequate when passing DC; contacts are often rated for a much higher voltage for AC than for DC.

My junction boxes consist of standard household breaker panels mounted on the garage wall. The panels are easy to mount. The entrances are designed for standard cable clamps and the wire terminations are all ready-made solid screw terminals.

For my 12- and 48-V systems (separate panels), I purchased and installed an extra "neutral bar" in the panel, isolated from the box, so I can have half the system on each bar and bring the two halves together electrically only at the inverter. I then tied all the negative leads to the "neutral bar" and all the positive leads to the circuit breakers. The backplane then became a bus bar and I could connect or disconnect solar panels, batteries, the inverter, the monitor, and DC loads (a 12-V automotive-style outlet) with the mere flip of a circuit breaker handle.

All this electrical hardware can be readily serviced by any electrician using standard tools. This fact is simultaneously an advantage and a drawback. I already have all the tools needed. However, an uninformed electrician getting into that box could have tragic consequences.

I have taken precautions at two levels. First, every such panel has a clearly printed label on the door as to what's inside. But if someone removes the cover and a different person starts working inside, there's a second level of notice. In standard house wiring, the neutral always uses white wire, and the "hot" phase wires use black, red, or rarely some other bright color, but never white. In a normal AC panel, all the white wires go to the neutral bar and all the black or colored wires go to the breakers. For exactly this reason I chose to use black for all the negative wires, which is consistent


with standard DC wiring. Therefore, all the black wires go to the neutral bar and all the white wires to the breakers. If one look doesn't set off alarm bells in an electrician's head, something's wrong!

By using 100-A rated breakers to connect the batteries and inverter and 15-A breakers for the solar panels, I can switch on or off individual panels (on the 12-V system) or strings (on the 48-V system), as well as batteries, and the inverter. Standard breaker panels enable you to pop out one breaker, visibly guaranteeing that you're disconnected from the bus, and deal with one wire while you do your thing with the screwdriver. In the case of a solar panel wire, I prefer to work at night (no power from the panel). I can pop out the breaker and then do as I will with the wire, capping it as needed if I need to leave it that way.

For the batteries, I will turn off the breaker, disconnect at least one terminal at the battery, and then know that I can work as needed on that wiring. My final result on the 48-V system has 470 AH of 48-V battery, and those batteries will happily deliver many kiloamps. Picture the result of an unprotected 100-KW surge across your wrench or your wiring. All those circuit breakers give me great comfort.

Notice how the system voltage has played an integral part in all of the design. On the one hand, a higher system voltage requires more panels in series and raises more safety issues, but the system current is lower and the inverter can be more efficient than at a lower voltage. These issues tend to push more toward a higher system voltage on a larger system.

## WHAT'S NEXT?

Now that I have the panels up and collecting all that free energy and a place to store it, the dirty work is mostly done. In Part 2 of this article series, I will discuss options for battery charge control, conversion to readily usable forms, and ways to conveniently share the load between the utility and solar power. 

## ABOUT THE AUTHOR

Engineering consultant Steve Hendrix (stevehx@hxengineering.com), CFIAI, lives with his wife Kathy and five children near Cuyahoga Valley National Park in Ohio and runs Hx Engineering, LLC. Steve has designed software and hardware for a wide range of projects: Space Station Freedom, vibration analysis, IEEE-488 bus hardware, biomedical sensing, industrial instrumentation, phototherapy equipment, and safety gear for facility personnel.



# DesignSpark Tips & Tricks

## Day #12: Using Buses

By **Neil Gruending**  
(Canada)

Today let's look at how to use buses in your DesignSpark schematics. Last time we explored multiple page schematics and the cross probe tool using a more complex example in DesignSpark. We'll continue that theme today and learn how to use buses in our schematics.

### Schematic Bus Types

Buses are very common in schematics because they are a great way to graphically group like signals together and make the design easier to understand. They are usually drawn as a wide

line with named net connections branching to and from the bus as in **Figure 1**. The wide blue line is the bus and the lines that end on the bus are the connected nets named Data0 to Data5. Let's take a closer look at buses in DesignSpark and the two different types that it supports: open and closed.

An open bus doesn't keep track of connections—it's just a line in the schematic. Open bus net connectivity is specified using the net names that automatically appear when you connect it.

**Figure 2** shows an example of an open bus. When I drew the connections from the resistors to the bus, DesignSpark automatically put the small 45° angle at the end of the connection and displayed the net name. If you delete the bus, all of the resistor connections will be deleted as well. You can change the net name by right-clicking on the net and selecting "Change Net."

A closed bus is more intelligent because it acts like a container for all of the bus nets and only those nets are allowed to connect to it. You change an open bus (the default) to a closed bus by right-clicking on it and opening its properties window. There you can click on the Add Net button to open the Add Net to Bus window shown in **Figure 3**. A list of all of the nets in the design will be presented (N0001 to N0004 in this case) that you can then choose to add to the bus. If you want to add a new group of net names that end in a number like an address bus, you can use the New Nets option. For example, if you wanted a 16-bit address bus you could set the name to A0 to A15, which are created when you click the "Add To List" button.

### Open vs Closed Buses

So which kind of bus should you use? It's a hard question to answer because

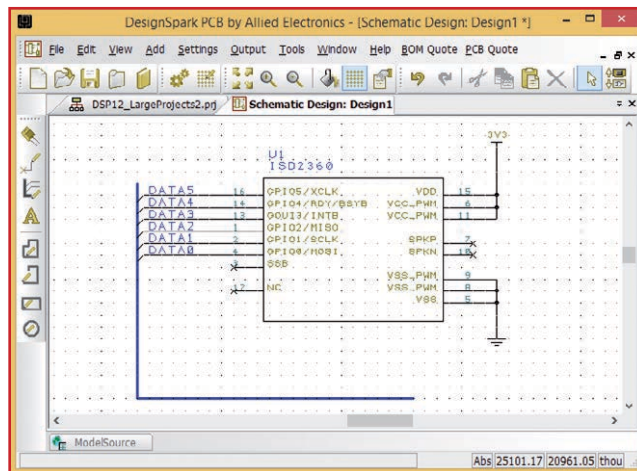


Figure 1.  
Bus Example Design.

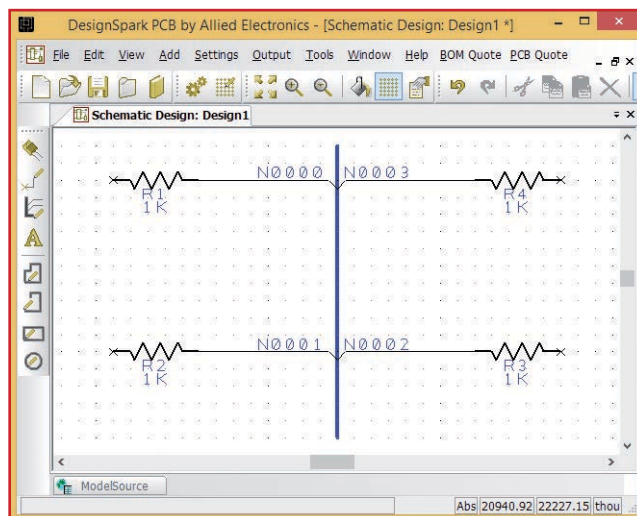


Figure 2.  
Open bus.

I think both types have their merits. I generally prefer to use closed buses for large buses such as a data bus and open buses for short small buses such as a SPI bus. But how DesignSpark treats them might also affect your decision.

Closed buses can be copied and they will keep all of the net associations, which can be useful for large buses. But be careful when copying between different schematic pages because the nets in a bus follow the same rules as regular schematic nets. So make sure you name them something different than the default Nxxxx. Also, when you paste the bus, DesignSpark will ask you if you want to merge the nets because it will detect that the net names already exist in the design. Normally, you would let DesignSpark merge the nets together.

Another advantage of closed buses in DesignSpark is that they can be probed using the new cross probe tool. When you probe a closed bus the cross probe tool will cycle through all of closed buses in the design that have the same nets associated with them. This is a great way to see where a bus is used in a design.

Open buses can also be copied, but when you paste them all you get is their shape. You can also probe an open bus too, except that you're limited to probing each net individually instead of the entire bus.

## Working with Buses

When you add a bus to a schematic, it draws a solid line that you can then make connections to. DesignSpark automatically creates a 45° bus connection point to the bus when you add bus connections, and you can control the size of the connection point line by using the Bus Defaults window in the Settings menu. It's also possible to change the direction of the connection line by clicking on it and using the F key (flip command).

One downside of buses is that it's very easy to miss a connection and make single node nets. DesignSpark doesn't have a schematic net check tool and the PCB design rule check (DRC) tool can't check for single node nets. But it is possible to manually check your design. The easiest way to do this is to open the DesignSpark project file and run the "Generic Netlist" report from the Output & Reports menu. It's important to do this

from the project file because if you run it from a schematic page the report will only include the nets from that page. **Figure 4** shows a report for a design with a closed bus with nets D0 to D3 and an open bus with A0 to A1, which are connected to several resistors R1 to R7. The report makes it easy to see that nets D2, D3, and A1 only have one connection. It's also easy to see that net D0 has three connections, which may also indicate an error depending on the design.

## Conclusion

DesignSpark's open and closed buses are easy to add to schematics and each has their strong points. The DesignSpark online help has even more details about how it implements buses and how to use them. Next time, we'll look at some of the more advanced DesignSpark PCB features.

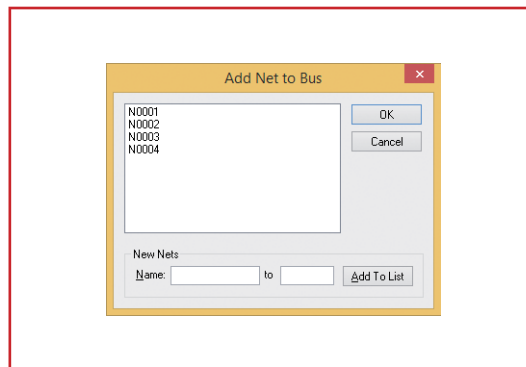


Figure 3. Add net to bus window.

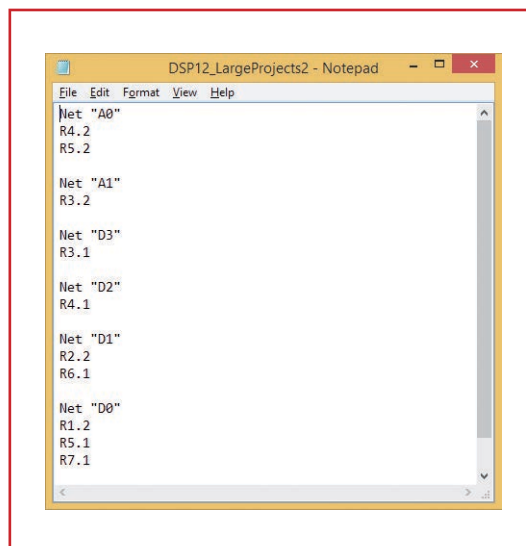


Figure 4. Netlist report.



# Tunnel Diodes

## Weird Component #7

By Neil Gruending  
(Canada)

I was looking for some inspiration around my lab this month and my Tektronix oscilloscopes caught my eye. I've always admired the older models and how they achieved such high performance with discrete components. One part that's always intrigued me was how Tektronix used tunnel diodes in some of their oscilloscope trigger circuits. They aren't common anymore but some new parts are still available if you look hard enough.

So why did Tektronix use tunnel diodes in oscilloscope trigger circuits? Probably because tunnel diodes switch quickly and can also act as a flip flop. For example, a Tektronix 465 oscilloscope uses a couple of tunnel diodes to create a trigger pulse for the oscilloscope logic circuitry. One diode is used to arm the trigger pulse and the other is used to fire it [1]. The diodes are necessary to allow the circuit to have defined operating states for the trigger logic.

**Figure 1** shows a tunnel diode in a simplified bias circuit. Tunnel diodes are made from a PN silicon junction with a very narrow depletion zone to

maximize the tunnel effect. This means that the diode will start to conduct forward current with very little bias voltage. The forward current will increase and eventually peak ( $i_1$ ) and then reach a minimum ( $i_2$ ) as the bias voltage is increased like in **Figure 2**. This behavior gives the tunnel diode its negative resistance characteristic and lets it be used for an amplifier or oscillator. The peak current value is one the key characteristics when specifying them.

Let's look at **Figure 1** again. The 10-mA diode is biased with 9.9 mA of current, which is right at the peak diode current when  $V_{TRIGGER}$  is high impedance. The diode is triggered into its minimum current state when  $V_{TRIGGER}$  injects enough current into it to push it past the peak current.

This triggering characteristic is also how you can test a tunnel diode. If you measure the resistance with a DMM and it reads a high value then the diode is defective. A low resistance reading means that it could be okay, but to know for sure you need to verify that the diode will switch between its peak and minimum current values. In some critical circuits you may also have to verify that the values haven't drifted over time, which is a common issue. If the values have changed, some circuit adjustment and calibration might be necessary.

Tunnel diodes are pretty interesting devices that met a real need in the 1960s and were heralded as the next big semiconductor discovery after the transistor. They're pretty hard to find now unless you are repairing vintage electronics. But if you're interested in learning more about tunnel diodes and their uses, check out some of the Tektronix oscilloscope enthusiast websites [2].




Figure 1. Simplified tunnel diode bias circuit.




Figure 2. Tunnel diode current characteristic.

### Web Links

- [1] <http://w140.com/mmm/tek-465-late.pdf>
- [2] [http://w140.com/tekwiki/wiki/Tunnel\\_diodes](http://w140.com/tekwiki/wiki/Tunnel_diodes)



# 2015 begins

# 9.9.2014

at Super Mobility Week  
*powered by CTIA*

**Super Mobility Week** is North America's largest forum for the mobile innovations that power your connected life. Thousands of mobile professionals and executives, 1,100+ exhibitors, as well as 1,000+ media and analysts from across the globe will gather in September. Register for this powerhouse, week-long technology event in Las Vegas to experience unrivaled networking opportunities and share leading-edge trends, discoveries and knowledge.

KEYNOTES

**Mobile Video**  
featuring a panel hosted by  
**Dan Mead**  
*President & CEO, Verizon*  
*CTIA Chairman*

**Connected Car**  
featuring a panel hosted by  
**Ralph de la Vega**  
*President & CEO, AT&T Mobility*

**Privacy**  
featuring a panel hosted by  
**Dan Hesse**  
*CEO, Sprint Corporation*

**Policy**  
**Tom Wheeler**  
*Chairman, FCC*

**Enterprise**  
**Bill McDermott**  
*CEO, SAP*

PARTNERS

## INTELLIGENT NETWORKS & SERVICES

CTIA 2014  
THE Mobile Marketplace

4G  
alltel

4C WORLD

CCA  
Communications & Content Association  
and Network Operators

CDC

FierceWireless

informa  
telecom & media

M2M

TOWER & SMALL CELL  
SUMMIT

WI-FI  
ALLIANCE

Bio

Continua

Wireless  
Life Sciences  
Alliance

THE  
MONEY  
EVENT

SNACHA

NAHB  
National Association  
of Home Builders

PARKS  
ASSOCIATES

## HEALTH

## MONEY

## HOME

## AUTO

IAA AUTO  
SHOW

Connected Car  
Ego

## MEDIA

CNBC

DIGITAL  
GMA

iPhoneLife

## PRIVACY

CTIA  
The Wireless Association

iapp

## RETAIL

IFA

TWICE

WIRELESS  
REPAIR  
EXPO

## ENTERPRISE & BUSINESS IT

Mobile  
CON

ibj

Award  
Solutions

CompTIA

MOBILE  
ENTERPRISE

temia

Mobility supercharges everything. Live for **Super.**  
**September 9, 10 & 11, 2014** Sands Expo Las Vegas




SuperMobilityWeek.com

# Automotive Isolated Switching Power Supply

FEATURES

Software simulation of a switching circuit can help you troubleshoot your design before problems occur. This project uses a SPICE program to model a 12-to-5-V isolated switching power supply before constructing and testing an actual circuit.

By Tom Struzik (US)



Visit <http://bit.ly/1u4nTWb> for a web tutorial that accompanies this article.

Some time ago, I designed a Bluetooth audio adapter for my car. (Refer to my article "Wireless Data Exchange: Build a 2,700-lb. Bluetooth Headset," *Circuit Cellar* 240, 2010.) There was one problem with the design: supplying power to keep the phone charged and to power the Bluetooth adapter. To run both, I needed a clean, quiet, 5-V USB-compatible power supply. It needed to be capable of providing almost 2 A of peak current, most of which would be used for the smartphone. In addition, having an in-car, high-current USB power supply would be good for charging other devices (e.g., an iPhone or iPad). The problem was, the typical "cigarette lighter" car chargers were only 5-V/1-A devices. Thus, it appeared to be time for another project. How hard could it be to throw together a quick 10-W/5-V auto power supply with a couple of USB jacks?

I began by considering the simple solutions. A power supply built from a 5-V/2-A linear regulator was off the list without much thought. The power loss from dropping the car's 12 V down to 5 V was a quick showstopper. Next was an inexpensive, prebuilt 12-to-5-V/2-A switching converter. I hoped a switching converter would solve the power loss problem, but when I tried it, the problem was noise. Serious noise! A bit of troubleshooting quickly pointed to alternator whine, ground loops, and high-frequency noise from the switcher itself. That eliminated all the simple solutions from the list. Now what?

Some Internet searching yielded two main understandings. First, the ground loop between the car radio AUX<sub>IN</sub> and the Bluetooth headset circuit AUX<sub>OUT</sub> had to be eliminated. Two distinct power supplies tied to different ground points but connected



PHOTO 1

Here are the results of simulating the original 12-/5-V design. This graph shows the primary, secondary, and output voltages. This is a good place to start your own design. (Image based on LT3748\_TAO2.asc from Linear Technology's website.)

via the audio signal line simply would not work. Second, the switching power supply needed a “quieter” design. These combined requirements pointed to the solution of using an isolated switching power supply. This sounded like a good follow-up to my previous Bluetooth headset project.


My first assumption was that switching power supply design was a mystery and it left me wondering if I really could design my own switching power supply. I did another quick Internet search and found Linear Technology’s LTspice SPICE simulator program. The program is a nice graphical SPICE simulator with built-in models for Linear Technology’s switching regulators and downloadable and working switching power supply simulations including several designs for isolated power supplies. Now I had a place to start.

### INITIAL DESIGN

The first side effect of choosing the LTspice simulator was that the resulting circuit design would likely be based on one of Linear Technology’s switching regulators. Filtering Linear Technology’s parts list for the requirements of input/output voltages, isolation, and components that had design guides and readily available LTspice models took me almost directly to Linear Technology’s LT3748 switching controller and its associated design guide “LT3748 100-V Isolated Flyback Controller.”

The design guide was the definitive resource for building a switching power supply with the LT3748 controller. It outlined the overall process for building various power supplies with this chip. This included the theory of operation, formulas for selecting components, general circuit tuning, and a complete walkthrough of a 12-V<sub>IN</sub>, 5-V/2-A<sub>OUT</sub> design. There was even a prebuilt, downloadable LTSpice simulation for a 12-V<sub>IN</sub>, 5-V/2-A<sub>OUT</sub> switcher similar (but not identical) to the walkthrough example in the design guide. It sounded perfect.

I started the design work by running the 12-/5-V simulation “as-delivered.” As an ideal simulation, it produced a near ideal result, as expected (see **Photo 1**). However, very quickly the real world began to intrude on this simulation. The first issue was the transformer. Digi-Key was the only place to find the LT3748 in small quantities, so ideally I’d get the transformer from them too. Unfortunately, the simulation only specified a generic transformer model and the transformer detailed in the design guide example was not available at Digi-Key. Second, some changes would have to be made to the circuit if the final design was going to survive in the planned automotive environment. I had



**PHOTO 2**

This primary voltage and current has a 2:1 transformer configured for a 5-V output (a). Compare this with a 3:1 transformer (b) and a 4:1 transformer (c). As the turn ratio increases, the primary’s peak voltage goes up and the peak current goes down.

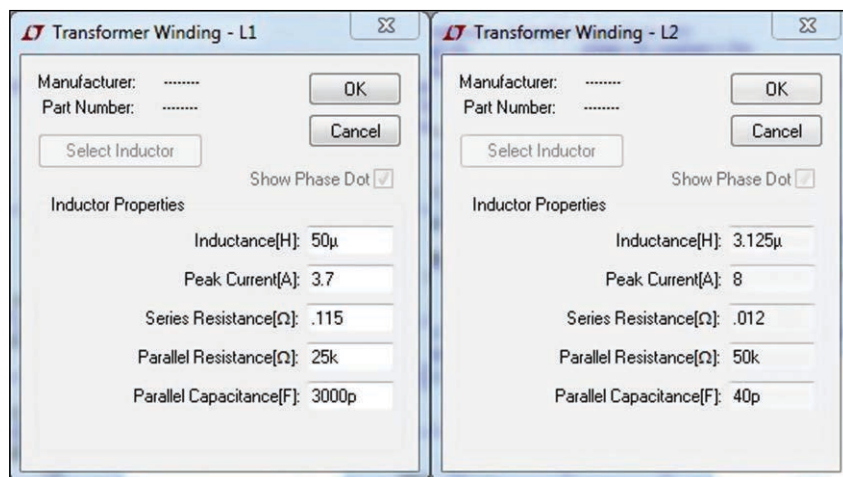


PHOTO 3

These are the typical inductor parameters available for modeling. The coupling factor between the primary and secondary coils, which appears as a SPICE parameter on the circuit drawing, is not shown. I set the coupling to 0.999 to get a more realistic model.

## Efficiency Report

Efficiency: 59.4%

Input: 16.8 W @ 14.5 V

Output: 10 V @ 5 V

Reference	I <sub>RMS</sub>	I <sub>PEAK</sub>	Dissipation
C1	0 mA	0 mA	0 mW
C2	211 mA	1,791 mA	0 mW
C3	1 mA	1 mA	0 mW
C4	0 mA	0 mA	0 mW
C5	145 mA	270 mA	3 mW
C6	1,705 mA	5,032 mA	9 mW
C7	1,233 mA	3,731 mA	2 mW
C8W	82 mA	155 mA	1 m
C9	1,705 mA	5,032 mA	9 mW
D1	4,005 mA	12,166 mA	1,201 mW
D2	0 mA	0 mA	1 mW
D3	0 mA	0 mA	-0 mW
D4	1,196 mA	1,827 mA	487 mW
Q1	1,735 mA	2,827 mA	1,653 mW
R1	0 mA	0 mA	8 µW
R2	0 mA	0 mA	2 mW
R3	0 mA	0 mA	52 µW
R4	0 mA	0 mA	0 µW
R5	0 mA	0 mA	169 µW
R6	0 mA	0 mA	23 µW
R7	1,699 mA	2,950 mA	36 mW
U1	338 mA	1,985 mA	642 mW
U2	4,005 mA	12,166 mA	210 mW
U3	2,080 mA	2,081m	2,567 mW

TABLE 1

This is a preliminary efficiency report for the circuit with the 5-V regulator. Note the poor efficiency: 20% is lost in the regulator itself, 16% is lost in the field-effect transistor (FET), and 12% is lost in the output diode. Needless to say, this particular iteration didn't stick around very long.

to dive into Linear Technology's design guide in more detail.

To tweak the circuit, it helped to understand the overall design. The design guide does a good job there. The key points were selecting the transformer, the switching field-effect transistor (FET), and the output diode. These components also had to be modeled rather than being represented by simple SPICE components. With this basic knowledge in hand, I knew I had to start with the replacement transformer selection.

Linear Technology's design guide lists (and Digi-Key sources) various Würth Electronics transformers, all with different turn ratios and ratings. For a given output requirement, a smaller turn ratio meant lower voltages and higher currents in the primary (see **Photo 2**). Thus the FET requirements were directly dependent on the transformer selection. Cost was also a driver. A 2:1 transformer cost around \$21, a 3:1 transformer cost about \$18, and a 4:1 transformer cost around \$15. The designs with the lower turn ratio/higher cost transformers also ended up with a higher current and potentially more expensive FET. A FET for a 2:1 transformer cost around \$8, while a FET for a 4:1 transformer cost around \$2. The end analysis made it appear that the lower turn ratio designs ended up being more expensive.

Based on specification datasheet and model availability within LTspice, I focused on utilizing a Würth Electronics 750311689 transformer and a Vishay Siliconix Si4490DY FET. The transformer had a 4:1 ratio, acceptable voltage/current ratings, and a 50-µH primary. The FET was rated at 200 V with a 2.8-A continuous current. As a first pass, this combination at least appeared to be in the ballpark, but time would tell as I moved through simulation to actual circuit construction.

After I included the selected components within the LTspice simulator, I was able to prototype various designs without building anything physical. However, the more complicated the simulation, the longer it took to run. Eight hours or more was not uncommon. So choose your simulation runs wisely.

Several LTspice simulations I experimented with are available on *Circuit Cellar's* FTP site so you can download and try them. If you find yourself scratching your head over why a particular change in the simulation fails to operate as expected, first check the voltages and currents within the circuit. LTspice doesn't fail like a real circuit does with heat and smoke; it fails by not converging on a solution, or worse it fails by providing a solution, but one that is nonsensical.

As a side note, Würth Electronics has

LTspice models for some of its transformers including the 750311689, which I had selected. However, for some unknown reason, I was unable to use Würth Electronic's transformer models to build an accurate simulation. But, because the company's transformer datasheets are fairly comprehensive, I was able to piece together a transformer model simply based on two coupled LTspice inductors (see **Photo 3**). In the long run, this ended up being a good move as it enabled me to experiment with the transformer model's internals to gain a better understanding of how the various transformer parameters impacted the overall switcher design.

One general design issue I continued to fight was ripple in the output. I initially had a standard 200- $\mu$ F electrolytic for the output capacitor, but the high current along with a "normal" equivalent series resistance (ESR) capacitor was contributing to a relatively strong output ripple in the 100-kHz range. Given the high peak currents in the secondary side, it was important to keep the output capacitor's ESR as low as possible. In fact, I quickly discovered it was better to have multiple low ESR capacitors in parallel rather than one large one. Modeling in the simulator quickly confirmed this. It's also worth noting that the default capacitor model in LTspice comes up with 0 ESR (i.e., an ideal capacitor, rather than a real one). Since my ultimate intent was to produce a real circuit, it was important to go through all the critical components and adjust the component models to be as close to "real life" as possible.

Every new circuit eventually needs its first smoke test. Even though the simulation doesn't actually generate smoke and heat, I found the LTspice efficiency report was helpful in producing simulated smoke to test a new circuit design (see **Table 1**).

The efficiency report calculates the amount of power loss (i.e., heat) in each component.

Apple Device Charging Control		
Maximum charging current	D-	D+
500 mA	unconnected	unconnected
1,000 mA	2.8 VDC	2 VDC
2,000 mA	2 VDC	2.8 VDC
Android Device Charging Control		
500 mA (a.k.a., on USB)	unconnected	unconnected
2,000 mA (a.k.a., on AC)	shorted	

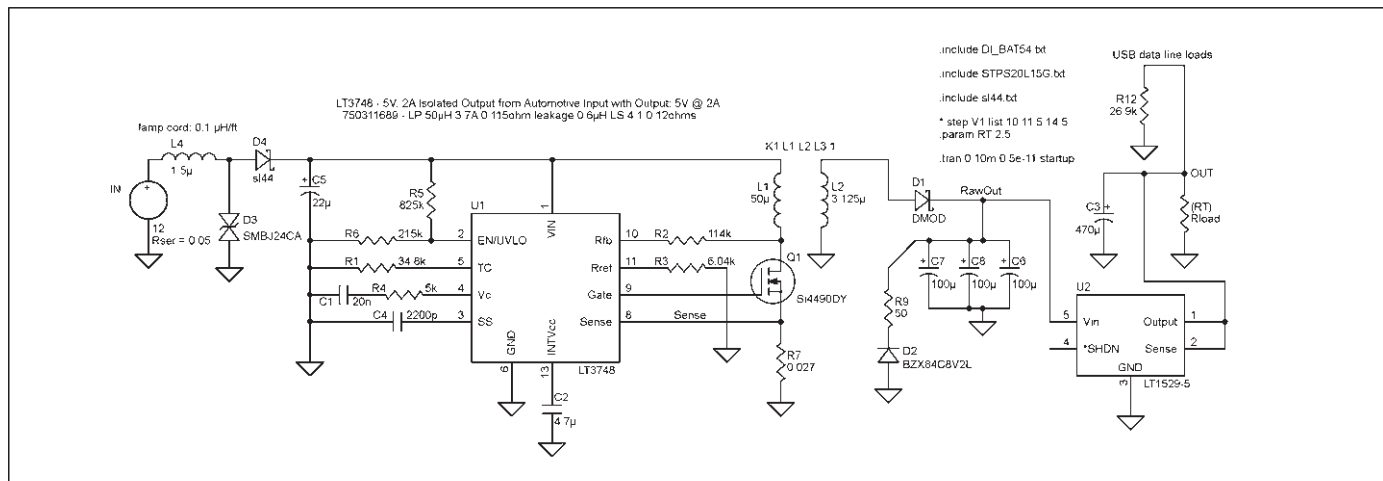
After some tweaking, I thought I had a good initial design, but the efficiency report still showed issues: a lot of power loss. The main offender at this point was the output diode and, sure enough, the design guide pointed out that diode's critical nature. If this had been a real circuit rather than a simulation, I probably would have had a lot of smoke!

Examining the design guide further I found that the output stage diode had to be fast and have a low forward loss. Choosing a better diode was a daunting task, until I limited the search to only those diodes for which the manufacturer provided SPICE models. That was a short list and a major constraint of using this method to model a real circuit.

After achieving what I thought was a simulation reasonably close to real life, a few extra changes were still needed. Specifically, changes to the input/output portions of the circuit. On the input side, the circuit had to survive in an automotive environment. I found an application note on Littelfuse's website that details some of the major requirements for the automotive environment. Those included extreme requirements (e.g., surviving a 24-V jump start and a 125-V, 200-ms transient in a "load dump" situation when the car battery is disconnected while being charged by



**TABLE 2**  
USB charging specifications are shown.

**FIGURE 1**  
The initial circuit design with the 5-V low dropout (LDO) regulator as represented in Linear Technology's LTspice is shown. Note the embedded SPICE commands, some of which were not used in this specific run. With variables and step commands, you can vary element values during a simulation run.



**PHOTO 4**

This is my initial test circuit board, post hacking. A Zener diode is shown in the upper right, a multi-turn trimmer for feedback resistor is in the center, a snubber capacitor and “stacked” surface-mount design (SMD) resistors are on the center left, USB D+/D- voltage adjust trimmers are on top center, and a “test point” is shown in the far lower left. If you’re looking for the 5-V low dropout (LDO) regulator, it’s on the underside of the board in this design.

**PHOTO 5**

The effects of the snubber on the primary and output voltages are shown. Without the snubber, a 20-V spike on top of a 30-V square wave appears **(a)**. With the snubber, only a 10-V spike and fewer cycles of ringing are shown **(b)**. The results of this are shown in **(c)** and **(d)**, which are the A/C component of the output voltage. They display 200-mV peaks without the snubber and 50-mV peaks with the snubber, along with a reduced number of ringing cycles.

the alternator). Then on the output side, I needed USB ports that would charge devices at higher than the “standard” 500-mA level.

I covered the input requirements by simply including a bidirectional transient voltage suppression (TVS) diode to limit positive/negative transients. Then I used a fuse to limit the maximum current through the diode. I also added a reverse-protection diode and a low ESR input capacitor to filter some of the other inherent environmental noise. The only problem with all this was that the switcher’s input voltage was further reduced. This meant tweaking the feedback resistor value in the circuit and rerunning a bunch of the simulations to ensure things were stable.

The output requirements ended up with their own unique problems. If all I had wanted was a meager 500-ms USB charger, this would have been simple. However, because I wanted a high-current charger, mysteries and problems appeared.

It seemed the device manufactures could not agree on a high-current USB charging standard. Thus to charge most Android phones, the USB D+/D- lines must be shorted together. However, to charge an iPhone or iPad, the D+/D- lines have to be held to very specific voltages to convince an Apple device to charge at a high current. Worse, Apple didn’t appear to want to publish its methods, so I was left searching the Internet. **Table 2** shows my search results.

Testing also showed that to really achieve some of the higher charging currents, the output voltage had to be at the very upper end of the USB standard range (around 5.2 V). The end result was the requirement for the switcher circuit to have two USB ports, one configured for Android charging and a second configured for Apple charging. Actually, based on the 2007 USB Battery Charging Specifications, Android followed the specifications and Apple did not.

There was one final item on the output requirements list. The switching circuit always had to be presented with some load even if nothing was plugged into the output. The Linear Technology design guide quotes the minimum load at 2% of the rated output. Rather than simply putting a resistor or LED on the output, which would always drain current, the design guide suggested using a Zener diode, which only switches on in the unloaded state. This works because the output voltage floats upward without any load, so the Zener diode only draws current in an elevated voltage, no-load state.

After cleaning up the noise in the circuit as much as I could, I was still dissatisfied with the overall output noise level predicted by the simulator. My final solution was to give up. Well, not quite. My penultimate final solution

## ABOUT THE AUTHOR

Tom Struzik (tpstruzik@earthlink.net) lives in Houston, TX. He has a BSEE from Purdue University. Tom is an IT systems architect at a Fortune-100 chemical company’s engineering organization. He enjoys writing software and building hardware projects at home to “keep his hands dirty.”

was to place a 5-V low dropout (LDO) regulator on the tail of the circuit. Of course this entailed boosting the output voltage a bit more to give the regulator something to work with, but overall it seemed like a good compromise. Thus the switcher does most of the work stepping the highly variable 11.5-to-14-V input down to

# \$51<sup>For 3</sup> PCBs

## FREE Layout Software!

## FREE Schematic Software!



- 01 DOWNLOAD our free CAD software
- 02 DESIGN your two or four layer PC board
- 03 SEND us your design with just a click
- 04 RECEIVE top quality boards in just days

expresspcb.com




PHOTO 6

I tweaked this simulation to exhibit ringing similar to the real-world circuit. Compare this simulated primary output voltage with the actual measured voltage in Photo 5a. It's not exact, but it's close enough.

a more reasonable 5.8 V. The regulator then only has to handle a relatively stable 5.8-VDC input and a small voltage drop. Yes, adding the regulator cut the circuit's overall efficiency, but at least the model now showed a nice quiet output voltage (see **Figure 1**).




PHOTO 7

This is the primary output voltage in the simulation complete with modeled ringing and a snubber. Compare this with the actual measured voltage in Photo 5b.

## REAL-WORLD APPLICATION

**Figure 2** shows the finalized circuit design. With this done, I moved from modeling to the real world. My circuit board design is not difficult. The only exception was that the LT3748 had a fairly small 16-MSOP footprint. It's certainly hand-solderable, but it's a bit "fidgety" to achieve.

Also, here are a few suggestions if you design your own board. You will have to tweak some of the resistor and capacitor values, so add a few additional footprints in parallel and series to some of the more "tweak-worthy" components. Leave enough space around the components so you can get a soldering iron in to remove and replace the components you intent to tweak. Finally, add a few test points so you can attach clips to measure critical voltages. Unfortunately, my initial circuit board design did not include these suggestions, so I had to hack these in after the fact (see **Photo 4**).

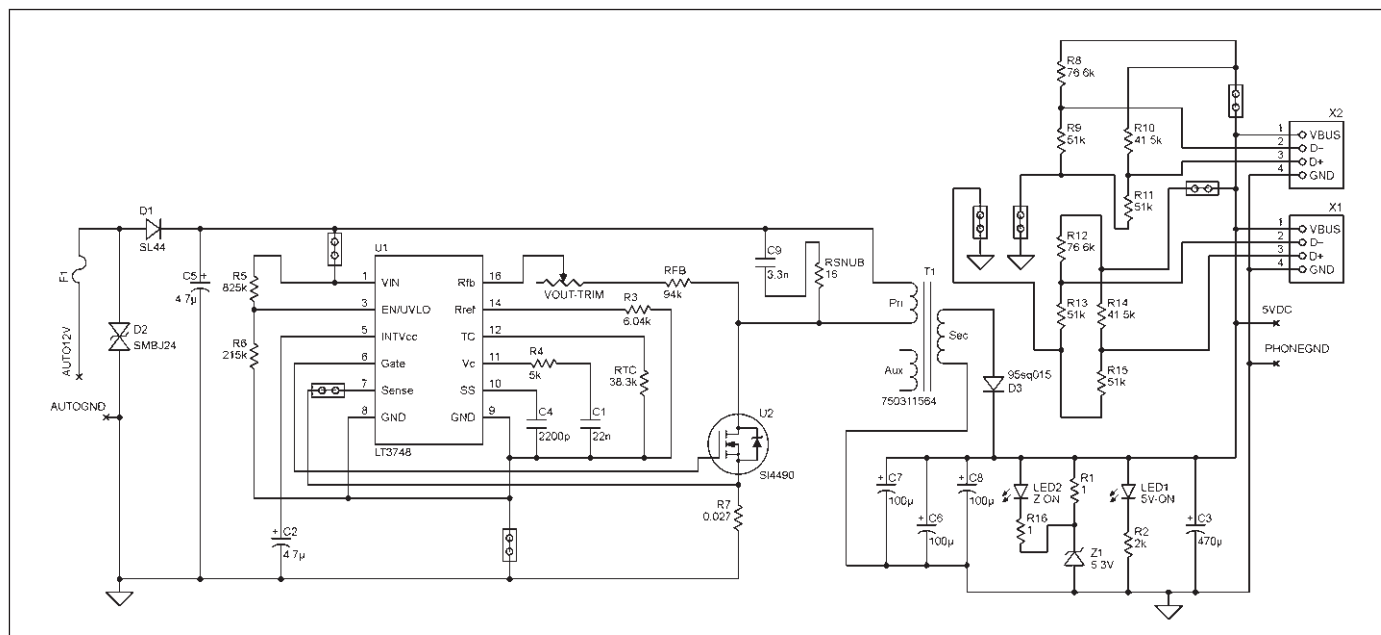
After getting the new circuit board in hand and soldering in the components, I was surprised to find that it appeared to work on the first attempt, 12  $V_{IN}$  in, 5  $V_{OUT}$ . However, the real world decided to intrude again. There was more noise in the output than the model predicted. Worse, my archenemy, heat, showed up again. The regulator, the output diode, and even the transformer were running near the "too hot to touch" point. I suppose this should have been expected, given that the LTspice simulation is more "ideal" than the real world and all those real-world losses are the tough parts to model in the simulation.

It turned out the heat and the noise were related. First there was unexpected noise in the output due to "ringing" when the FET switched. (I wondered why there was no ringing in the model?) That noise forced the output voltage set point even higher to keep the regulator from dropping out. Then the 12-V power supply I used for testing could not keep up with the switcher's peak demands.

Therefore, with all the required output voltage headroom and the reduced input voltage, the current through the transformer was substantially higher than expected. In fact, it was high enough to force the transformer into saturation. It was obvious this was the wrong approach, so I removed the regulator from the simulation and the board and then tweaked the feedback resistor to get the output voltage down to 5 V. The operating temperatures returned to "warm" but the output was still more noisy than acceptable.

I pulled out the design guide again to find a solution to the ringing. I thought this must be a fairly common problem as the design guide not only had the solution, but included





**FIGURE 2**


This is the final circuit design with the 5-V regulator removed and the snubber inserted. I left a trimmer as part of the feedback resistor to enable final tweaking of the output voltage. Note the odd “jumpers” in some of the ground and power lines are there to break the signal so the auto-router will route the majority of the high current lines with wide traces, then the other side of the “jumper” can be a thin, low-current trace. Doing this, if you’re careful, the board will auto-route 100% of the traces even with the very wide traces.

a complete work process for identifying the necessary component values. All it required was a simple RC snubber across the transformer primary coil.

Determining the approximate component values was simple enough and only required an oscilloscope. The process was basically to add enough capacitance across the primary to approximately double the ringing period. Then I needed to calculate the resistance based on formulas in the design guide and it was done. And yes, this did substantially lower the output noise. Also, the output voltage went up 0.25 VDC when the snubber was added, so the overall circuit gain could be reduced, which further lowered the noise (see **Photo 5**).

One question remained: Why didn’t the ringing show up in the simulation? After some experimentation with LTspice, I got some of the ringing into the simulation by decreasing the transformer coupling and increasing the primary inductor parallel capacitance. The couple factor seemed to drive the magnitude of the ringing while the capacitance drove the period. This got the simulation substantially closer to the real-world circuit (see **Photo 6** and **Photo 7**).

At this point, close was good enough and so it was on to the true test: How does it sound? I moved to the car, plugged the power supply into the car’s cigarette outlet, connected the phone’s USB to the switcher and the headphone jack to the car’s AUX<sub>IN</sub>, and started listening. It had a nice sound.

There was no alternator hum and no ignition noise. It looks like I have a winner. For more information about this and other projects, visit [www.jenrathbun.com/electronics](http://www.jenrathbun.com/electronics). 

### PROJECT FILES



[circuitcellar.com/ccmaterials](http://circuitcellar.com/ccmaterials)

### RESOURCES

Linear Technology Corp., “LT3748 100-V Isolated Flyback Controller,” 2010.

Littelfuse, Inc., “Suppression of Transients in an Automotive Environment,” Application Note AN9321.5, 1999.

Sixerdoodle Electronics, “LTspice Tutorial—Car USP Power Supply,” [www.jenrathbun.com/Electronics](http://www.jenrathbun.com/Electronics).

T. Struzik, “Wireless Data Exchange: Build a 2,700-lb. Bluetooth Headset,” *Circuit Cellar* 240, 2010.

Wikipedia, “USB: Charging Port and Accessory Charging Adapters.”

### SOURCES

**LTspice SPICE simulator and LT3748 controller**

Linear Technology Corp. | [www.linear.com](http://www.linear.com)

**Si4490DY Field-effect transistor (FET)**


Vishay Siliconix | [www.vishay.com](http://www.vishay.com)

**750311689 Transformer**

Würth Elektronik GmbH & Co. | [www.we-online.com](http://www.we-online.com)

# Build a Thermal Monitoring Network

FEATURES



You can use ANT+ wireless technology to build a management system that provides real-time temperature monitoring. The system described here features software that configures the network, provides temperature value information, detects alarm levels, and stores and exports data.

By Carlo Tauraso (Italy)

**I**t is often necessary to monitor the temperature of different rooms or equipment units. Data loggers have various limitations (e.g., high cost per detection point, the dedicated wiring, and the static addressing).

The current communications market offers a range of technologies based on open standards (e.g., Bluetooth, ZigBee, 802.11bgn, and RFID). However, in recent years the electronics market has focused on decreasing costs and energy consumption while requiring

the same performance levels. This trend has highlighted some open-standard limitations. For example, many developers in an open-standard software community place different requests on a protocol that are irrelevant to a specific application, which makes it difficult to optimize cost and performance in a common layout for the standard.

In 2010 Nordic Semiconductor introduced a new chipset that incorporates a proprietary open-access wireless technology called ANT+. This technology, managed by Dynastream Innovations, runs on the 2.4-GHz band and enables the realization of networks especially optimized for communication between devices with a low data rate and low latency.

I decided I could combine the need to measure temperature with some of ANT+ technology's interesting technical features. A network consists of a coordinator node connected via USB to a computer that uses RF to query a series of temperature probes.

I built a management system to measure within a 100-m radius in real time without the




FIGURE 1

The coordinator sends an addressed packet and node 0A replies.

need for wires. You can use my ThermoNET software to configure the network, view the temperature values, detect alarm levels, and store and export the data to other commercial programs. One of the project's interesting aspects is that the remote node firmware is loaded directly into the transceiver core's system-on-a-chip (SoC), which minimizes the need for external components. The project provides a good basis for the development of a low-cost temperature-monitoring system.

## ANT+ TECHNOLOGY AND PROTOCOL

Nordic Semiconductor, which is a leading manufacturer in this field, initially received a contract from Garmin (the well-known manufacturer of a professional GPS navigator system) for the construction of a series of products to be integrated into navigation devices and enable communication with other external devices (e.g., heart rate monitors, bike cadence sensors, etc.). The ANT+ technology Nordic Semiconductor introduced is positioned at an intermediate level between Bluetooth and ZigBee in terms of data transfer speed and drastically decreases the complexity and cost of production.

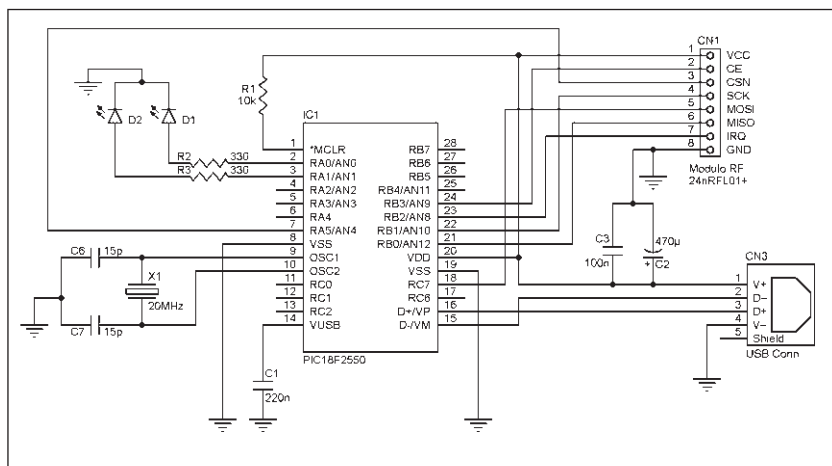
Nordic Semiconductor's design focused on energy saving. My design uses a star configuration with a coordinator node connected to the PC via USB and one or more remote nodes carrying a digital temperature probe. The remote nodes only communicate with the coordinator node through a master-slave relationship.

In a personal area network (PAN), using a dedicated channel to communicate is often imposed for each pair of devices. Although there are 125 available channels with 1-MHz separation in the 2.4-GHz ISM band, I preferred to use shared-channel configuration.

The ANT+ protocol can use up to five bytes as destination addresses. Using only two bytes, it is possible to communicate with up to 65,000 probes on a single channel, which overcomes the expandability limitations of many commercial data loggers.

The coordinator controls all transmissions. When it wants to query a particular node, it sends a packet with a payload specifying the remote node's address. All nodes receive the packet but it is only processed by the node whose address matches the one in the payload. Only that node can then use the channel to transfer the response. No other node can engage the transmission channel.

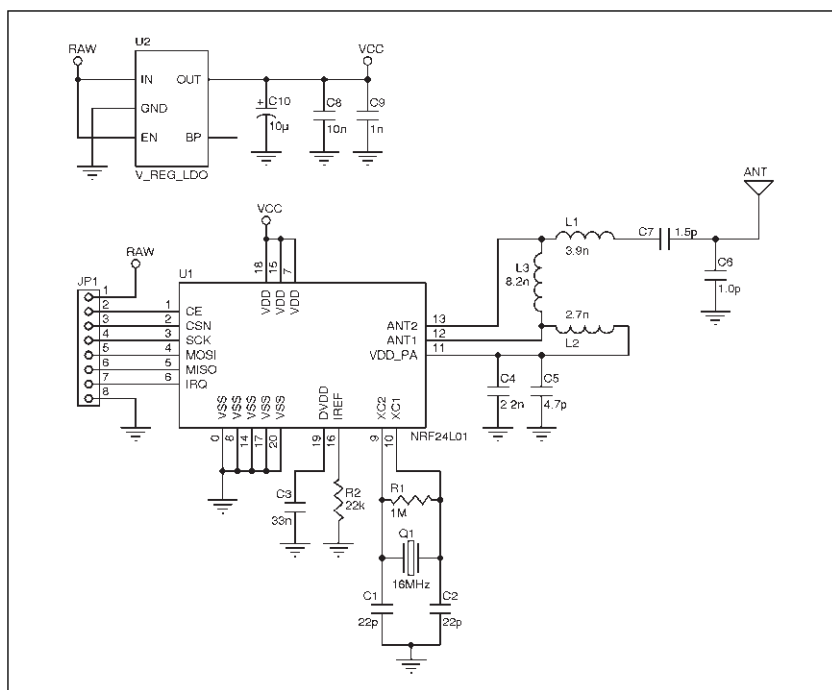
For example, suppose there is a coordinator with three remote nodes that have the addresses 000A, 000B, and 000C hex (see **Figure 1**). The coordinator sends a packet with the first two bytes 0A 00. The



**FIGURE 2** The RF module is directly connected to the Microchip Technology PIC18F2550 microcontroller to optimize firmware development.

least-significant byte is transmitted first. Only node 0A (highlighted in yellow) processes the message and has the opportunity to respond. The flow direction is indicated by arrows.

The protocol was designed by Dynastream Innovations, which is a subsidiary of Garmin. The protocol is very compact and requires minimal resources. The physical, network, and transport Open Systems Interconnection (OSI) model layers are encapsulated within the ANT chipset, so you only need to interface



**FIGURE 3** The RF module schematic is based almost exclusively on Nordic Semiconductor's nRF24L01 RF transceiver. Only some passive components and an external 3.3-V voltage regulator are needed. (Schematic information courtesy of SparkFun Electronics)




PHOTO 1

The final coordinator prototype is less than 4" long. The RF module is shown on the right.

a host microcontroller with the ANT module at the OSI application layer. ANT is channel-based, so each node can connect to other nodes via a dedicated channel. Normally there is a master node and one or more slave nodes. The master and the slaves must have common knowledge of the channel configuration (network key, channel ID, RF) to maintain communication sessions.

I used a shared channel in my design, so slave nodes share a single channel to communicate with the coordinator node. The microcontroller application and the ANT module communicate through a simple serial interface.

Every message has a fixed basic format implemented in little Endian. It begins with a preamble byte followed by an address field (three to five bytes), a 9-bit packet control field, and a payload field (zero to 32 bytes), and ends with a checksum. The preamble is a bit sequence used to synchronize the receiver's demodulator to the incoming bitstream. The address field is the receiver's address.




PHOTO 2

The ceramic antenna saves space although the performance is slightly lower than with an ordinary antenna.

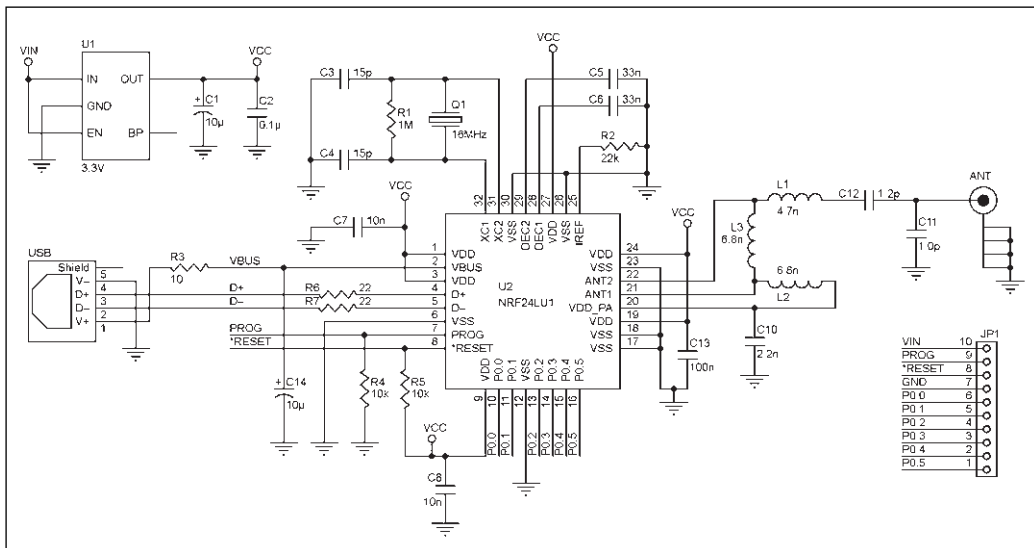
My design uses 2-byte-long addressing. The packet control field specifies the payload length, the packet identification, and an ACK flag. It is used only with the ShockBurst feature for automatic retransmission of loose packets. The payload is user-defined. My design only uses two bytes for temperature value. The cyclic redundancy check (CRC) is calculated as a polynomial function calculated over the address, packet control, and payload. The packet is assembled directly by the ANT module. You only need to assign the right values to the right internal registers. If you want to know more, the protocol details and everything you need to develop with ANT are available on Dynastream Innovations's website.

## THE COORDINATOR

The combined circuitry shown in **Figure 2** and **Figure 3** forms the coordinator shown in **Photo 1**. **Figure 2** shows the circuit, which consists of a Microchip Technology PIC18F2550 and a module based on Nordic Semiconductor's nRF24L01 RF transceiver. This is the least expensive chip currently distributed by Nordic Semiconductor. It integrates all the necessary components for the transmission and reception in the 2.4-GHz ISM band and interfaces with the outside through a SPI port. The transceiver enables the management of three independent TX/RX 32-byte first in, first out (FIFO) buffers. It works at 3.3 V but all of the interface lines are 5-V-tolerant.

The resistor on pin 1 (MCLR) is required to execute the firmware in the microcontroller. It functions as an active-low reset line. Two LEDs (green and red) indicate the circuit's status. Resistors R1 and R2 limit the current that flows through them. I used the MSSP module inside the PIC18F2550 to interface with the RF module. I used pins RB0, RB1, RC7, and RA5 as serial input serial data input (SDI), serial clock (SCK), serial data output (SDO), and slave select (SS). The last two pins of the RF module CE and IRQ are used respectively to activate/deactivate the transmission or reception and handle any interrupt from the module.

**Figure 3** shows the board with the unregulated 5-V input coming in from the USB port on the board from **Figure 2**. **Photo 2** shows the SparkFun Electronics WRL-00691 board I used for the prototype. **Figure 3** is the board's schematic. The oscillator circuit consists of a 16-MHz crystal with two capacitors C1 and C2. The input pin VIN of the 3.3-V voltage regulator is connected to the USB bus' 5-V pin. The capacitor C3 has decoupling function on pin DVDD. For space reasons, I used the model with the integrated ceramic antenna. The circuit has an impedance network adapter with capacitors and inductors (L1, L2, L3, C6, C7) to enable the connection of 50-Ω



**FIGURE 4**  
The impedance network adapter with capacitors and inductors is shown on the right of the breakout board schematic.

omnidirectional antennas such as those used in the home routers. Adapting the impedance is critical to avoid losing power. I added a small plastic box with two holes on the top for the LEDs (see **Photo 1**).

**THE END DEVICE**

The combined circuitry shown in **Figure 4** and **Figure 5** forms the coordinator

shown in **Photo 3**. A Nordic Semiconductor nRF24LU1+ SoC is at the heart of the circuit (see **Figure 4**). The SoC integrates an 8-bit microcontroller that is compatible with an Intel 8051 microcontroller’s instruction set, a USB controller, and finally 32 KB of flash memory. The chip has six I/O programmable lines and works with 3.3 V. I used a SparkFun Electronics WRL-09034

**The Convenient All-in-One Solution**  
for Custom-Designed Front Panels & Enclosures

**FREE Software**

**You design it**  
to your specifications using our FREE CAD software, Front Panel Designer

**We machine it**  
and ship to you a professionally finished product, no minimum quantity required

ONLY \$90.24 with custom logo engraving

- Cost effective prototypes and production runs with no setup charges
- Powder-coated and anodized finishes in various colors
- Select from aluminum, acrylic or provide your own material
- Standard lead time in 5 days or express manufacturing in 3 or 1 days

**FRONT PANEL EXPRESS**  
FrontPanelExpress.com  
1(800)FPE-9060

**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   BBB




PHOTO 3

The final end-device prototype shows the reset button, the probe, and the high-brightness LED. I placed a lithium-ion (Li-ion) battery in the space between the two boards.

breakout board for the prototype (see **Photo 4**). I added my circuit with a Maxim Integrated DS18B20 sensor, two jumpers for reprogramming and power, a high brightness LED for status, and a lithium-ion (Li-ion) battery.

I arranged a strip of connection above the circuit that includes the necessary lines to a small breadboard containing the various components (see **Photo 3**). The temperature sensor on the end device, the DS18B20, is powered using Dallas Semiconductor's one-wire parasite power mode. Power is supplied through a 2.2-k $\Omega$  resistor on the data line DQ. The circuit was completed by the bridge between the VDD and GND (see **Figure 5**).




PHOTO 4

At the bottom of the breakout board, 10 circular pads easily interface the transceiver with other external circuits.

The battery has been inserted into the space between the two circuits to minimize the overall space. The jumper CN4 enables you to turn the device on or off. The jumper CN2 normally opened was placed between the line P0.3 and the ground. Through the pull-up resistor R1, the line is always high except when the jumper is shorted. This will start the Firmware Update mode.

The reset line already has a pull-up resistor as seen in the schematic; therefore, I added one switch that is used to force a system hardware reset. The P0.5 pin is connected to a high-brightness LED (D1). In this way, you can check the probe's status even at distances greater than the transmitter's actual range. I also added an inexpensive 50- $\Omega$  omnidirectional antenna.

## THE FIRMWARE

I developed firmware[A] that is loaded on the PIC18F2550 to manage the coordinator and the communication between the microcontroller and USB and firmware[B] that is loaded on the nRF24LU1+ core to manage the DS18B20 and communicate between the end device and the coordinator. For the firmware[A], I used the C18 language and the Microchip Technology USB framework. The PIC device works as a human interface device (HID). I used this type of development because the HID drivers are built into all modern Windows systems so you will not have to do any type of installation. The device will be directly recognized by the OS. For PC communication, I used a simple request/reply protocol with eight bytes (see **Table 1**).

For interfacing with the RF module SPI port I used a function to read/write one byte to the SPI PIC buffer. **Table 2** lists the commands available from Nordic Semiconductor.

## ABOUT THE AUTHOR

Carlo Tauraso (carlotauraso@gmail.com) studied computer engineering at the University of Trieste in Italy and wrote his first Assembler code in the 1980s for the Sinclair Research ZX Spectrum. He is currently a senior software engineer who does firmware development on network devices and various types of microinterfaces for a range of European companies. Several of Carlo's articles and programming courses about Microchip Technology PIC microcontrollers (USB-PIC, CAN bus PIC, SD CARD, C18) have been published in Italy, France, and Spain. Carlo enjoys playing with radio scanners and homemade metal detectors.

The coordinator receives a command via USB from a computer that is running the monitoring software. It sends a command packet to the remote node, receives the response containing the temperature data, and at the end transmits the received data to the PC.

Firmware[B], which is loaded into the firmware of the end device nRF24LU1, receives commands from the coordinator. It uses the DS18B20 probe to measure the temperature and sends the collected data to the coordinator. I used Code Blocks as the IDE, which is a free, open-source product. The compiler is the small device C compiler (SDCC) distributed under the GNU General Public License.

First the firmware writes a series of internal registers that enable you to configure the RF module. The transceiver has five internal pipes, each with its own physical address. Each pipe works in parallel on the same RF physical channel. I only used Pipe 0. Each node will have a unique 2-byte RX address. I used a 1-mW transmission power level and a 250-Kbps data rate on Channel 2. To avoid delays in the TX/RX cycle, the temperature measurement is carried out immediately after the sending to the coordinator. The 12 bits sampling involves a delay of at least 750 ms; therefore, I decided to do so in the dead time between one transmission and the next. All communications with the DS18B20 must follow the sequence: Initialization, ROM command, Function command.

The initialization sequence consists of a reset pulse sent on the data line followed by a presence pulse from the probe. The nRF24LU1 pulls the data line (P0.4) low for at least 480  $\mu$ s, then releases it and goes into Receive mode. When the line is released, the pull-up turns the line high. When the probe detects this rising edge, it waits 60  $\mu$ s and then transmits a presence pulse by pulling the data line low for 60 to 240  $\mu$ s. The Command ROM phase is used when multiple probes are physically connected on the same bus. I used the Skip ROM command (value CCh) to skip this step since the probe is unique and is identified by the wireless network address assigned to the transceiver.

The function commands for this design are ConvertT (value 44h) to measure the actual temperature and Read Scratchpad (value BEh) to read nine bytes from the probe memory. The firmware must read all nine bytes and save only those corresponding to the temperature's most-significant bit (MSB) and least-significant bit (LSB). TX and RX operations are managed through two interrupt signals made available in the microcontroller configuration register (00h). The receive operation is ended as soon as the bit register RX\_DR (Rx Data Ready) is equal to 1.




FIGURE 5

In this schematic, the probe's GND and VCC pins are connected based on the particular power mode. Resistor R2 ensures sufficient current for the probe.

The transmission operation is ended as soon as the bit register TX\_DS (TX Data Sent) is equal to 1. The firmware cleans the registry after each operation. If the transmission is initiated and the TX\_DS is not active within a certain time interval, the packet is considered lost. Similarly, if you do not have a transmission after the response is received within a certain time, the remote node may not be working properly.

### THE SOFTWARE

The first program (ATU.exe) enables you to read and write the nRF24LU1+ internal flash memory on the remote node. I used C#, Microsoft's Visual Studio 2012, and NET framework 4.5. To update the remote node firmware, simply connect it to a PC via the micro-USB port and run the provided software. According to the manufacturer's specifications, the flash memory after production is preprogrammed with a small piece of code (2 KB) that works as a bootloader. It is recorded in the high portion of the internal flash memory. If you analyze the preloaded code, you realize that the reset

a)		
Byte	Name	Description
1	ID_CMD	Identifies the command the microcontroller must perform (e.g., 20h indicates a request for reading the temperature probe with the address specified in bytes 2, 3, 4, 5, 6)
2-6	IND_SONDA	The probe address to be queried
7-8		Not used
b)		
Byte	Name	Description
1	ID_RSP	Identifies the response to the command from the PC (e.g., 28h is the response identifier for the 20h command)
2-6	IND_SONDA	The probe address
7-8	TEMP	Temperature value in DS18B20 format

TABLE 1

A request/reply protocol identifies the microcontroller's commands and the PC's responses. a—This table shows the packet from the PC to the microcontroller in detail. b—This table shows the packet from the microcontroller to the PC.

TABLE 2

These basic commands are needed to interface with a Nordic Semiconductor nRF24L01 chip.

Command Name	ID Command (2 Byte)	Byte Number	Description
R_REGISTER	000A AAAA	From 1 to 5	Read register with address AAAA (5-bit register map address)
W_REGISTER	001A AAAA	From 1 to 5	Write register with address AAAA
R_RX_PAYLOAD	0110 0001 (61h)	From 1 to 32	Read RX-Payload. Payload is deleted from buffer after it is read.
W_TX_PAYLOAD	1010 0000 (A0h)	From 1 to 32	Write TX-Payload
FLUSH_TX	1110 0001 (E1h)	0	Flush TX buffer
FLUSH_RX	1110 0010 (E2h)	0	Flush RX buffer

vector (initial three bytes) points to the 7800h location. So as soon as you power the module, 2 KB of code are transferred into the static random-access memory (SRAM) and executed from there.

The bootloader contains a small command interpreter. The flash memory is divided into 64 pages of 512 bytes. Each page is divided into eight blocks of 64 bytes. The writing and reading are carried out for individual blocks. When you connect the node to the computer for the first time, you must install the drivers provided. After that, information about the manufacturer, firmware version, and the vendor ID/product ID will appear on the form.

To write new firmware in the flash memory, just left-click on the Update button. This opens a dialog box that enables you to choose the binary file to load. Be careful: if you write a file larger than 30 KB, you overwrite the bootloader code and lose the ability to reprogram the node.

The program also enables you to configure the remote node's network address. Each probe has two address bytes that are set by simply clicking in the left checkbox Assign Address and entering the relative value in the address field. In this way, when the firmware is loaded into the flash memory chip (Update button), the program writes the address area.

The second program (TNET.exe) enables you to configure the network, monitor the temperature in real time, and save the data in a file. So, connect the coordinator node to a free USB port. If the device is recognized by the OS, the green LED illuminates. At this point you can start the program. A window will appear (see **Photo 5**).

To insert a new node, type in the address and a description and configure a threshold level for the alarm. Only the first two fields are required. If there is something wrong, left-click on the Clear button. The fields will reset to their initialization values. Once you have entered various fields, left-click on the Add button and the node will be shown in the grid below. It's important to remember that all the nodes you insert are temporary, saving the final configuration will only take place after a left-click on the Save button. You can finalize the structure when you are satisfied with the result.

Even if a node has already been added to the grid it can be changed simply by left-clicking on the field you want to edit and pressing F2. For the alarm threshold, you can also use the buttons on the right side of the field. If you need to change the address, you have to delete the node and reinsert it. To delete a node from the grid, select the row and press the Delete key. Cancellation occurs only if you left-click on the Save button. You can reload the last saved network configuration at any

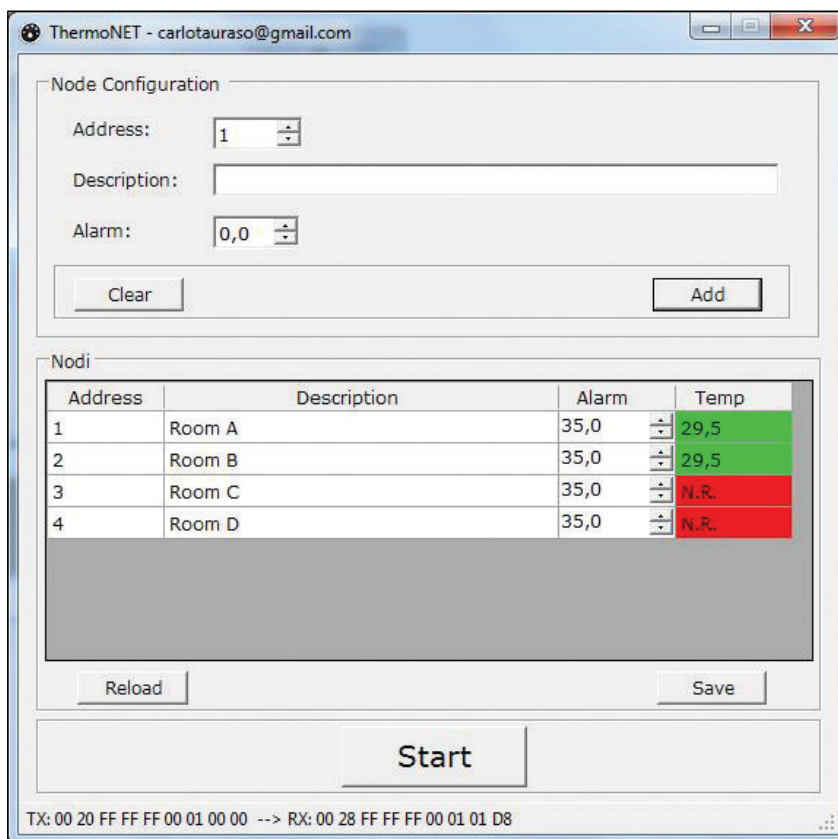


PHOTO 5

This screenshot shows four nodes being monitored, two of which are off.



time by clicking the Reload button. By left-clicking on the big Start button, the PC begins to record temperature values. It sequentially examines each node inserted in the grid. On the controller at each TX/RX cycle you will see the flashing red LED.


The grid reports in the Temp column the temperature values detected by the probes. If the measured value is higher than the threshold, the field turns red and you will hear a beep. The warning is repeated every poll until the value returns below the threshold. If the measured value is less than the threshold value, the field turns green.

There are two possible error levels. The first error is not received (N.R.), in which the transceiver does not respond because it may be out of range or have a flat battery. The other error is probe fault (P.F.), in which the transceiver responds correctly but cannot communicate with the probe, or it may be faulty or disconnected.

Pay attention to the fact that if you query a probe for the first time after it has been

powered, it will respond with the initialization temperature value (85°C). You can left click on the big Stop button to stop monitoring. In the directory where you save the TNET.exe you will see a file (.csv extension) with all values recorded.

## INEXPENSIVE WIRELESS TEMPERATURE CONTROL

This design represents a starting point for the construction of an inexpensive wireless temperature control system. There are many application fields. You can also use this design to better understand ANT+ networks. You can add more end devices simply by increasing the address length from two up to five bytes. You can also monitor other environmental parameters. If you need more information fields, you can set the payload length from two up to 32 bytes. The RF operating range is sufficient to cover your home area. For the moment, I hope this circuit can be useful for those who want to enter the world of ANT+. 

### PROJECT FILES



[circuitcellar.com/ccmaterials](http://circuitcellar.com/ccmaterials)

### RESOURCE

Dynastream Innovations, Inc., "ANT Basics: Welcome to the World of Ant," [www.thisisant.com](http://www.thisisant.com).

### SOURCES

#### 8051 Microcontroller

Intel Corp. | [www.intel.com](http://www.intel.com)

#### DS18B20 Sensor

Maxim Integrated | [www.maximintegrated.com](http://www.maximintegrated.com)

#### PIC18F2550 Microcontroller

Microchip Technology, Inc. | [www.microchip.com](http://www.microchip.com)

#### Visual Studio 2012 IDE

Microsoft Corp. | [www.microsoft.com](http://www.microsoft.com)

#### nRF24L01 RF Transceiver and nRF24LU1+ system-on-a-chip

Nordic Semiconductor | [www.nordicsemi.com](http://www.nordicsemi.com)

#### WRL-00691 and WRL-09034 Breakout boards

SparkFun Electronics | [www.sparkfun.com](http://www.sparkfun.com)

## RS485/422/232/TTL

### ASC24T

RS232<=>RS485 ATE Converter



\$45.00 board only

### IBS485HV

5 Port Isolated

RS485 Repeater  
\$349.00



Enclosures, Cables,  
Power Supplies and Other  
Accessories



### INTRODUCING THE SMFCOMX!

- Converters
- Repeaters
- Multi-Repeaters
- Hubs
- Fiber Optics
- Isolators
- Extended Distance Units
- Serial to Digital I/O
- Large Multi-Drop Networks
- Custom Units & Smart Units
- Industrial, 3.0 KV Isolation



Call the RS485 Wizards  
513-874-4796  
[www.rs485.com](http://www.rs485.com)

## GREEN COMPUTING

# Power Management in Mobile Devices

COLUMNS

Mobile computing systems now include several processor cores and more advanced graphics processing units (GPUs). These hardware advances combined with increased performance demands of mobile applications have motivated designers to find new power management strategies.

*By Ayse K. Coskun (US)*


**P**ower management of embedded systems has been a well-studied topic because embedded systems are traditionally severely limited in their power, thermal, and energy budgets. Despite the fact that today's embedded processors are significantly more power-efficient compared to prior processors, there is still a strong motivation for building new power management strategies. This motivation stems from the fact that the industry is now placing a larger number of processor cores, more advanced graphics processing units (GPUs), larger capacity memories, and other circuitry in embedded devices. In tandem, user performance demands from embedded systems continue to grow.

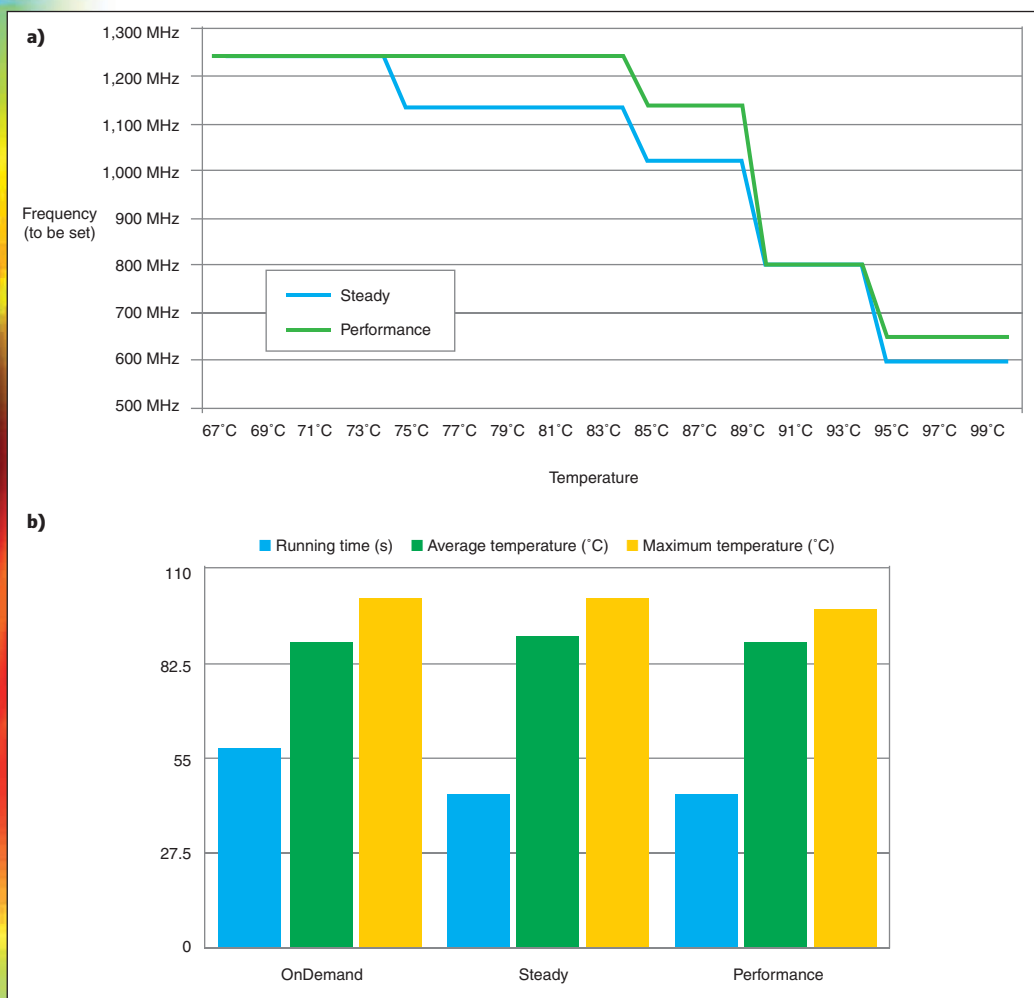
Consider mobile devices as examples. Users perform complex multimedia processing, modeling, document processing, or even artificial intelligence tasks (e.g., learning and recognition) on their mobile phones and tablets. Some common applications are gaming, social networking, web browsing, and various forms of interactions with cloud resources. This article discusses how the evolving hardware and software in mobile systems changes the way to design efficient power management mechanisms.

## POWER VS TEMPERATURE

Getting higher performance out of a processor typically implies higher power densities. This is especially a critical concern as the process technology shrinks and more functionality (i.e., power) is packed within a smaller chip area. Higher power densities create higher temperatures. In desktops, servers, or larger-scale systems, higher temperatures cause cooling costs to grow. In mobile systems, where cost-wise or area-wise it is not feasible to place sophisticated cooling solutions, higher temperatures place severe design and runtime constraints.

"Throttling" is one control knob available in a majority of recent processors for handling high temperatures. Throttling is essentially reducing a processor's voltage and frequency level (i.e., dynamic voltage and frequency scaling, or DVFS) or stopping the processor clock upon reaching a thermal threshold. As a larger number of computationally intensive tasks are executed on mobile systems, the processors may run in throttling mode for a larger portion of their active processing time. Throttling is a safety mechanism; however, frequent throttling is undesired as it may indicate performance loss.





**FIGURE 1**

**a**—This graph shows the temperature-dependent frequency levels set by each of the two policies. If the runtime temperature reaches a specific value on the X axis, the policy sets the corresponding frequency value on the Y axis. The Performance policy maintains the highest frequency until 85°C, while the Steady policy gradually reduces the frequency for temperatures above 75°C. **b**—The running times and temperatures when running the AnTuTu benchmark with the OnDemand, Performance, and Steady policies on the single-board mobile computer are shown.

In my research lab, my students and I investigated various DVFS policies that tune the voltage and frequency level without hitting the throttling threshold for mobile development platforms.


We ran DVFS experiments on an Inforce Computing IFC6410 single-board computer (SBC). The board includes a Qualcomm Snapdragon 600 system-on-a-chip (SoC) with four cores. We ran the AnTuTu Benchmark V. 4.3.2 on this platform.

**Figure 1a** demonstrates two policies for setting the DVFS level upon reaching a certain temperature level. We used the CPUfreq utility to implement these policies. CPUfreq “userspace” governors enable the user to implement arbitrary DVFS policies. In these two policies, we set the frequency level based on the temperature sensor readings. Temperature sensor data can be collected in our system by periodically reading a file (e.g., /sys/class/thermal/thermal\_zone0/temp). The Performance policy we implemented sets the CPU frequency at 1.242 GHz until it reaches 85°C to maintain higher performance. Then it sets lower frequency levels if the CPU reaches

higher temperatures. The Steady policy is less aggressive (i.e., it starts decreasing the frequency level gradually as the temperature increases to 75°C and higher).

**Figure 1b** compares the running time, average temperature, and maximum temperature when running the AnTuTu benchmark for the Performance, Steady, and OnDemand policies. OnDemand is one of the existing CPUfreq governors. It sets the CPU DVFS level based on CPU usage to improve energy efficiency. The temperature difference among the three policies is rather small. However, the running time shows variations. The shortest execution time of 44.8 s is achieved with the Performance policy. The second fastest policy is Steady, which achieves a 45.2-s running time. The OnDemand policy results in a 58-s running time. This indicates that policies such as OnDemand have considerable headroom to improve performance without hitting critical thermal thresholds. Note that always setting the frequency at the highest level statically (i.e., at 1.242 GHz) causes the board to eventually hit its thermal threshold and turn off.

These results indicate that there are



**FIGURE 2**

**a**—These are the running times of the AnTuTu benchmark at different frequency settings and with various numbers of active cores. My students and I compared setting a static frequency level all across the application with the dynamic OnDemand policy. **b**—The energy-delay product (EDP) for the same experiment is shown. Lower EDP is more desired.

opportunities to improve the application performance compared to existing governors, if desired. However, one should design policies that avoid or minimize temperature-induced throttling. One way of maintaining the temperature at desired levels is to dynamically track temperature through the available on-chip sensors as discussed earlier. Overall, I believe temperature awareness will be a central part of power and performance management on future high-performance embedded systems, as opposed to building thermal awareness solely to avoid reliability problems.

### DVFS VS THREAD PACKING

Future mobile systems will increasingly employ multi-core processors to improve their performance per watt. This increase in the number of cores introduces another control knob that can be used for power management: thread packing. I discussed thread packing as a power capping knob for multi-threaded applications running on servers in my article, “Application-Aware Power Capping” (*Circuit Cellar* 280, 2013). Now, I will discuss its potential benefits in the mobile systems domain.

Thread packing refers to placing the threads of an application (or multiple applications running on a system) onto a subset of the available cores. For example, on a four-core system, all active threads can be packed onto two of the cores and the remaining two cores can be turned off. Thread packing may reduce performance; however, it will also reduce power consumption and temperature. Therefore, it may be desirable as a power management knob.

In my research lab, we explored turning off one or more of the cores in a mobile system and running the AnTuTu benchmark with different DVFS settings and with various numbers of cores. In this experiment, we used a 8974-based Qualcomm Snapdragon mobile development platform (MDP) smartphone equipped with a Snapdragon Krait 400 four-core processor (**Photo 1**). We changed the DVFS setting through the CPUfreq userspace governor as before. Userspace governor also enables you to turn off one or more specific cores. We used this feature to explore thread packing. Note that thread packing can also be performed by setting thread affinities in the OS kernel as I discussed in my “Application-Aware Power

Capping" article.

**Figure 2a** shows the running times of the AnTuTu benchmark at different DVFS settings using one, two, three, or four cores. The general trend in the figure is that, as expected, the running time drops as the processor is run at faster frequency levels and a larger number of cores are turned on. Note that this MDP's processor has a much higher maximum frequency level available compared to the SBC used earlier (i.e., 2.15 vs 1.242 GHz).

While the general trend is as expected, we saw that the shortest runtime is achieved at the 1.49-GHz and three-core setting. Upon exploring the reason of this counterintuitive result, we observed that the processor was in throttling mode in some phases of the application when running at the highest DVFS setting with four cores. The three-core setting provided a slightly better tradeoff of temperature and performance and achieved a shorter running time in return.

**Figure 2b** demonstrates the energy-delay product (EDP) results for the same experiment. We obtained the EDP by multiplying the average power consumed while running the application with the square of the running time (power x running time<sup>2</sup>).



PHOTO 1

Ayse and her research lab team used a Qualcomm Snapdragon mobile development platform (MDP) to test various dynamic voltage and frequency scaling (DVFS) settings. (Photo courtesy of Qualcomm)

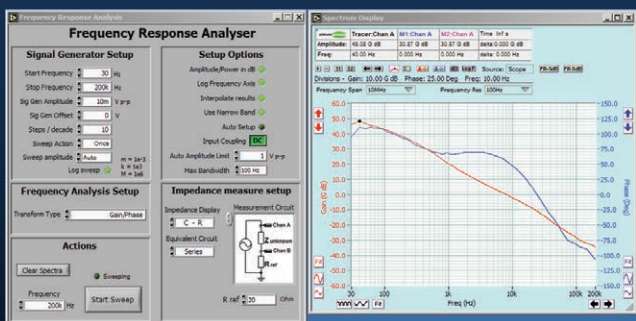
COLUMNS

## Frequency Response Analysis

Use the Cleverscope FRA panel to easily auto plot Gain/Phase, Impedance, Capacitance or Inductance vs Frequency. Display the Gain and Phase Margin.

Easy As, with Cleverscope.

See our FRA tutorial video to show you how to verify your operating power supply or amplifier design. Check the impedance of your DC buses. Verify magnetics you have wound. 80 dB dynamic range! 0 - 65 MHz isolated Sig Gen.



Streaming 100 G samples to disk ♦ Frequency Response Analysis ♦ Protocol Analysis ♦ Symbolic Math ♦ Matlab Interface ♦ 80 dB dynamic range ♦ 100 MHz Bandwidth ♦ Tracking Zoom ♦ 0-65MHz isolated sig gen ♦ Video Tutorials



**CS328A-FRA**  
**14 Bit MSO**



[www.cleverscope.com](http://www.cleverscope.com)

THE ORIGINAL SINCE 1994

# PCB-POOL

Beta LAYOUT

## PCB Prototypes & Medium Volume

Only at PCB-POOL®:

### 3D-data for EAGLE

**FREE!**

3D-PDF & Step file


3D mechanical check

3D-preview of loaded PCB

For information, please call toll free at 888-977-7443

**25 YEARS Beta**  
LAYOUT  
create : electronics

[www.pcb-pool.com/brd-to-3D](http://www.pcb-pool.com/brd-to-3D)



**FIGURE 3**

These are the running times of the execution kernels within the AnTuTu benchmark at different static frequency levels. The results for each kernel are normalized with respect to running that kernel with OnDemand. All runs use all four of the available processor cores.

The EDP is a commonly used metric for jointly evaluating energy and performance. A lower EDP is more desirable. For the mobile system and application with which we experimented, either running with the highest DVFS level but with fewer cores or running at medium DVFS levels with more cores provided the best energy-performance trade-offs (i.e., they achieve the lowest EDP). EDP dramatically increases when the DVFS level and the number of cores are reduced.

Is it better to favor a higher DVFS level or a larger number of cores to improve performance per watt? The answer to this question strongly depends on the specific application and the system. For applications with higher degrees of software-level parallelism, more cores may provide better results. For applications that are

computationally intensive, a higher DVFS level usually improves performance as the cores process instructions at a faster rate. In either case, our results show that one should simultaneously take performance, power, and temperature into account to find the most desirable operating point.

The AnTuTu benchmark includes various execution kernels. In another experiment, we looked into how each of these kernels responds to different DVFS settings.

**Figure 3** shows the running times of each kernel within AnTuTu when executed at a static frequency setting when running with four cores. The running times provided are in comparison to using the OnDemand governor. The CPU kernel consists of a CPU-intensive integer and floating-point loads. In the chart, RAM stresses the random-access memory, and the GPU kernel highly utilizes the GPU on the MDP with 2-D and 3-D graphics tasks. The I/O kernel performs a large amount of I/O operations, and UX mainly consists of user-experience intensive loads that are multitasking and Dalvik-oriented.

This experiment shows that performance of UX, RAM, and I/O are more sensitive to operating at lower frequencies. This is somewhat counterintuitive as one would typically expect the CPU-intensive loads to get a larger performance hit at lower frequencies compared to RAM or I/O. I believe this result points out two things. First, in mobile systems, performance of I/O and RAM operations are more closely coupled with the CPU performance compared to large-scale systems. Second, application-specific behavior strongly impacts the performance, and thus should be taken into account during power management.



circuitcellar.com/ccmaterials

[2] M. B. Taylor, "Is Dark Silicon Useful? Harnessing the Four Horsemen of the Coming Dark Silicon Apocalypse," in Proceedings of the Design Automation Conference (DAC), 2012.

#### RESOURCES

A. K. Coskun, "Application-Aware Power Capping," *Circuit Cellar* 280, 2013

#### SOURCES

##### IFC6410 Single-board computer

Inforce Computing, Inc. | [www.inforcecomputing.com](http://www.inforcecomputing.com)

##### Snapdragon 600 system-on-a-chip, 8974 mobile development platform, and Krait 400 four-core processor

Qualcomm, Inc. | [www.qualcomm.com](http://www.qualcomm.com)

#### REFERENCES

[1] I. Paul, S. Manne, M. Arora, W. L. Bircher, and S. Yalamanchili, "Cooperative Boosting: Needy vs Greedy Power Management," in Proceedings of the International Symposium on Computer Architecture (ISCA), 2013.

## CPU AND GPU COUPLING

Another emerging aspect of mobile system power management is the potential coupling between the CPU and the GPU. Coupling refers to runtime management decisions taken on the CPU that affect the performance, power, and thermal dynamics of the GPU and vice versa.

For example, given a total power budget, running the CPU at its highest DVFS level could leave less room for the GPU to run at its highest capacity. Depending on the specific application running on the target system, it could be beneficial to give a larger power budget to the CPU or the GPU.

This coupling effect is expected to grow for “fusion” systems that place a CPU and a GPU on the same chip. Recent work by Advanced Micro Devices (AMD) shows that performance coupling between the CPU and the GPU should be taken into account so as to maximize the performance under a total system power limitation.<sup>[1]</sup>

For example, reducing the DVFS level on the CPU may enable the GPU to run at a higher speed and improve performance for a GPU-intensive load. Reducing the CPU frequency too much, on the other hand, may cause stalls and degrade performance as the GPU and CPU often exchange data.

Another source of coupling is related to temperature. Running the CPU at a high temperature may heat the GPU as well (and vice versa) because of their close proximity in fusion systems. As a result, temperature-dependent leakage power may increase on the GPU or the CPU (or both), reduce the available power budget, and cause performance loss.

What are the implications of such coupling effects on mobile systems? First, as mobile systems typically run quite a few GPU-intensive workloads and are tightly space constrained, fusion systems may become desirable as they offer easier communication between the CPU and the GPU. CPU-GPU interactions on the chip will then add another dimension to the power management optimization. Second, even for systems with CPU and GPUs on separate chips, performance and power coupling effects still exist even though thermal coupling would likely be less significant.

Such interactions among the different components of an SoC motivate adopting a system-level strategy for designing power management policies. In a system-level approach, while each hardware component still has its own power-state controller (e.g., having DVFS and core turn-off capabilities in a CPU), the management decisions are guided considering the interactions among

various hardware and software components so desired performance and energy targets are met.

## THE AGE OF DARK SILICON

As the number of cores on a chip is growing even in the mobile systems, manufacturers face an additional power consumption challenge: dark silicon. Dark silicon essentially refers to the inability to simultaneously power on a substantial portion of a chip’s components in future technologies because of total power and thermal budget restrictions. In other words, manufacturers expect a large portion of a chip to be turned off or under-clocked at a given time in future chips.

Recent work explores how dark silicon can be leveraged to improve energy efficiency in various ways, including the use of specialized, application-specific circuitry.<sup>[2]</sup> Selecting the most-efficient hardware component to turn on (or to clock at a higher rate) or, more broadly, configuring the hardware depending on the application needs, is an especially beneficial feature for mobile systems because of the tight power budgets.

Future power management methods will need to be aware of the application-specific needs and characteristics, the physical interactions between power and temperature, the impact of temperature control mechanisms on performance, and the various forms of coupling among the system’s different components. Prior-generation embedded system power management policies were largely about making each hardware and software component individually more efficient. Next-generation strategies will be crossing the hardware and software boundaries and leveraging a full system-level view to meet aggressive performance demands under tight power and thermal budgets. [E](#)

*Author’s Note: Experimental results in this article were generated with the help of Boston University graduate students Linlin Zhu and Kushal Prasad.*



### ABOUT THE AUTHOR

Ayse K. Coskun (acoskun@bu.edu) is an assistant professor in the Electrical and Computer Engineering Department at Boston University. She received MS and PhD degrees in Computer Science and Engineering from the University of California, San Diego. Ayse’s research interests include temperature and energy management, 3-D stack architectures, computer architecture, and embedded systems. She worked at Sun Microsystems (now Oracle) in San Diego, CA, prior to her current position at BU. Ayse serves as an associate editor of the IEEE *Embedded Systems Letters*.

## THE CONSUMMATE ENGINEER

# WWVB Clock Revisited



COLUMNS

Atomic clocks located far from the National Institute of Standards and Technology (NIST) time signal radio station can produce weak synchronization signals. This article discusses the WWVB synchronization time standard, which was created to fix the problem of sporadic clock synchronizations.

By George Novacek (Canada)

In my article “Sine Wave Synthesizer” (*Circuit Cellar* 252, 2011) I vented frustration caused by my “atomic” clocks’ lack of synchronization. Although they are called “atomic,” they are regular tuned-fork resonator-controlled timepieces, synchronized via radio by a time standard located near Fort Collins, CO. The synchronizing data arrives as a 60-kHz amplitude-shift-key (ASK) modulated carrier transmitted by the National Institute of Standards and Technology (NIST) radio station with the call sign WWVB. Since 2012, NIST also transmits a phase-modulated (PM) signal. All my clocks are ASK, so I cannot comment on the PM receivers’ performance.

## “ATOMIC” CLOCKS

Due to the distance, reception of the WWVB signal where I live, in Eastern Ontario,


Canada, is unreliable. Resonator-controlled clocks (i.e., pretty much all battery-operated clocks) keep fairly accurate time. My “atomic” clocks without synchronization seem to be drifting significantly, but every now and then they do synchronize. Twice a year, when the daylight savings times change, the clocks do not gain or lose an hour for weeks, even months. When my impatience gets the better of me and I manually adjust the time, the clocks seem to get a good signal and jump an additional hour ahead or back. Subsequently, this calls for another manual adjustment. Worse, no two clocks respond at the same time.

I have a bunch of those clocks, although I no longer live by them. My retirement gift was a wall clock with only one hand showing days of the week; the hours can be estimated. For the most part, that’s all I need, but for an engineer, the “atomic” clocks present a challenge that cannot be ignored. It was suggested to me that the simplest solution may be to junk those clocks, but that would be an engineering cop-out. I solved the problem eventually, but it got a little bit more convoluted than I expected.

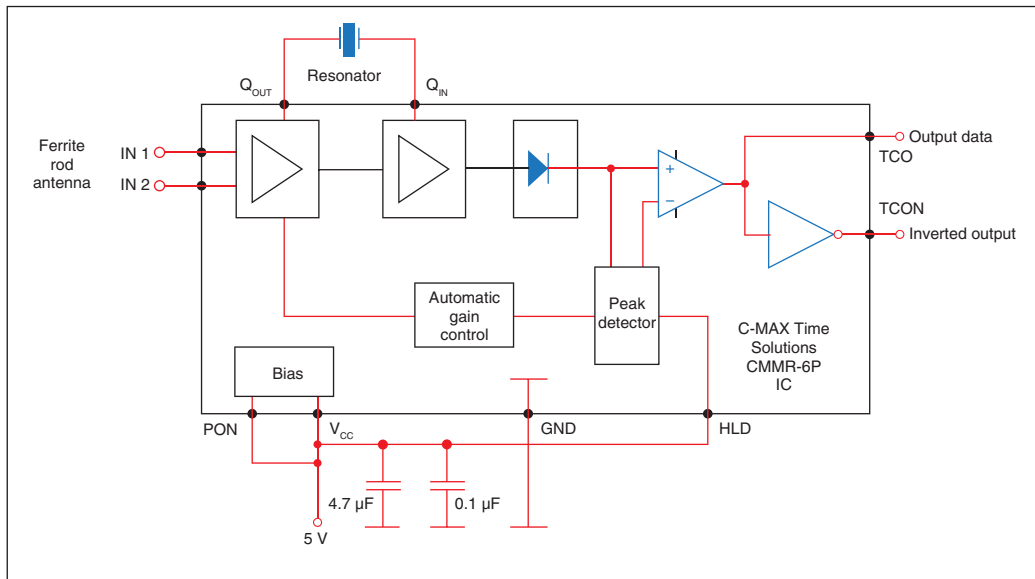
## SYNCHRONIZATION

The synchronization protocol is well defined in J. P. Lowe and K. C. Allen’s “Increasing the

FIGURE 1  
The WWVB bits definition is shown.







**FIGURE 2**

A C-MAX Time Solutions CMMR-6P IC receiver module provides the WWVB time signal reception.

Modulation Depth of the WWVB Time Code to Improve the Performance of Radio-Controlled Clocks,” paper. **Figure 1** shows the bits’ timing. Each data frame lasts 1 min and contains 60 1-s bits.

All my clocks came with 50-mm (1.97”) ferrite stick antennas. To increase the receivers’ sensitivity, I replaced those antennas with 100-mm (3.94”) ones. There was no perceptible improvement in the reception. I built a WWVB monitor to assess if there was a way to get a reasonably reliable signal by improving the clocks’ receivers. Later, I used the monitor to test my own transmissions. **Figure 2** and **Photo 1** show the monitor, which is built around C-MAX Time Solutions CMMR-6P monolithic receiver with a 100-mm ferrite stick antenna.

Even when I replaced the ferrite stick with a loop antenna with a preamplifier mounted on a tower, the reception remained unacceptably sporadic. Mostly, the WWVB signal stayed deeply buried in noise. That meant I had to generate my own synchronization.

I decided to utilize the GPS time standard I built some time ago to serve as a master clock for other projects. It generates a full time and date message once every second. I’ll provide more details about this later.

First, I focused on my repeater’s transmission antenna. After several tests I settled on a loop consisting of two turns of about 700’ of wire around my basement’s inside perimeter. The loop is made of several pieces of 22 AWG and 24 AWG wire. The wire is not critical, but all the splices should be soldered.


The loop has a 4.1-Ω DC resistance. With a 0.022-μF capacitor, it resonates at 60 kHz, exhibiting a 120-Ω impedance. The Q factor is quite low, about 9.6 with a 6.2-kHz bandwidth.

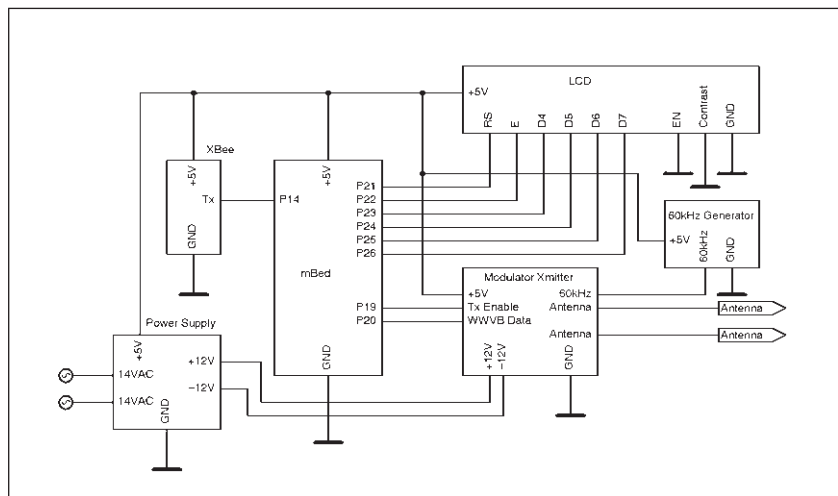
Its characteristics are affected by my basement’s duct work, wiring, water pipes, and so forth. I would not expect to obtain exactly the same characteristics elsewhere. However, other than bringing the antenna to resonance and adjusting its impedance matching, nothing else needs to change.

With about 9 mW of 60-kHz average power pumped into the antenna, all of my home’s clocks receive sufficient signal strength to synchronize. About 5’ outside the house, the signal is attenuated to the noise level. With

**PHOTO 1**

The WWVB signal reception monitor is built around a C-MAX Time Solutions CMMR-6P IC receiver.





**FIGURE 3**

The WWVB repeater station includes an LCD, antennas, a generator, an mbed LPC1768 microcontroller, and an Adafruit Industries XBee transceiver.

the nearest neighbor living 100' away, I don't need to worry about unwanted emissions.

## THE REPEATER

**Figure 3** shows the repeater's schematic. The GPS time standard is located in my office on the second floor because, even with an external GPS antenna, I couldn't get a reliable satellite fix in my basement workshop.

A long time ago, I was told in no uncertain terms to keep my wires where they cannot be seen. Since stringing and fishtailing wires through walls and ducts is not among my favorite pasttimes, I built the GPS standard with an XBee Series 1 transceiver. Its regular transmissions reach at least 100'. By

purchasing another XBee transceiver from Adafruit Industries, I can sync any device to it. Stringing the repeater's antenna in the basement was also the result of the "no visible wires" policy.

## WWVB SIGNAL GENERATOR

**Figure 3** shows a simple WWVB signal generator. The power supply provides  $\pm 12$  V for the modulator/transmitter analog circuits and 5 V for the microcontroller and logic circuits. The serial datastream received by the XBee transceiver feeds into an mbed LPC1768 microcontroller's UART. The time is displayed on the LCD and also converted into WWVB protocol, which feeds into the modulator/transmitter. I used the display to help me troubleshoot and as a time display in my workshop.

The "TX Enable" input enables me to limit the sync data transmission to certain times, currently between 4:00 AM and 5:00 AM. More frequent synchronization is not needed.

The 60-kHz carrier generator is a sine-wave emulator I described in my "Sine Wave Synthesizer" article (see **Figure 3**). I built my first repeater prototype with Arduino, but between decoding the incoming XBee serial data, validating it, updating the LCD, and encoding the WWVB data for transmission, it was unable to keep up. I already had an unused LPC1768, so it was an easy fix.

## CIRCUITRY

**Figure 4** shows the modulator/transmitter circuitry. It is heavily low-pass filtered. The unfiltered sine wave emulator's output wave has 21% total harmonic distortion. While the clocks synchronize even with the carrier being a square wave, transmitted signal purity has always been important to me. I needed three amplifiers. I used two Texas Instruments LM7332s, which are dual op-amps, so I utilized the extra amplifier as an additional filter. The result is a sine wave with 0.2% total harmonic distortion (THD) derived from a 60-kHz square wave.

The input R3C3 is a dual-purpose low-pass filter. First, it compensates for the op-amps' limited bandwidth in the Salen-Key filters that follow. (Refer to my article, "Filters, Filters Everywhere," *Circuit Cellar* 243, 2010.) It also attenuates the 60-kHz emulated sine wave to 380 mV RMS, which follows through all amplifiers with Gain=1 and represents the carrier's low amplitude. Two Salen-Key low-pass filters follow. The modulator U3B is AC coupled to avoid DC current through the antenna. When Q1 is turned on by the WWVB data, the stage's gain is increased by 17 dB for 100% modulation level.

I used Adafruit Industries' Ultimate



circuitcellar.com/ccmaterials

The National Institute of Standards and Technology (NIST), "WWVB Radio-Controlled Clocks," [www.nist.gov](http://www.nist.gov).

### SOURCES

Wireless XBee products and Ultimate GPS breakout board

Adafruit Industries | [www.adafruit.com](http://www.adafruit.com)

### CMMR-6P IC Receiver

C-MAX Time Solutions GmbH | [www.c-max-time.com](http://www.c-max-time.com)

### LPC1768 Microcontroller

mbed | <https://mbed.org>

### LM7332 Op-amp

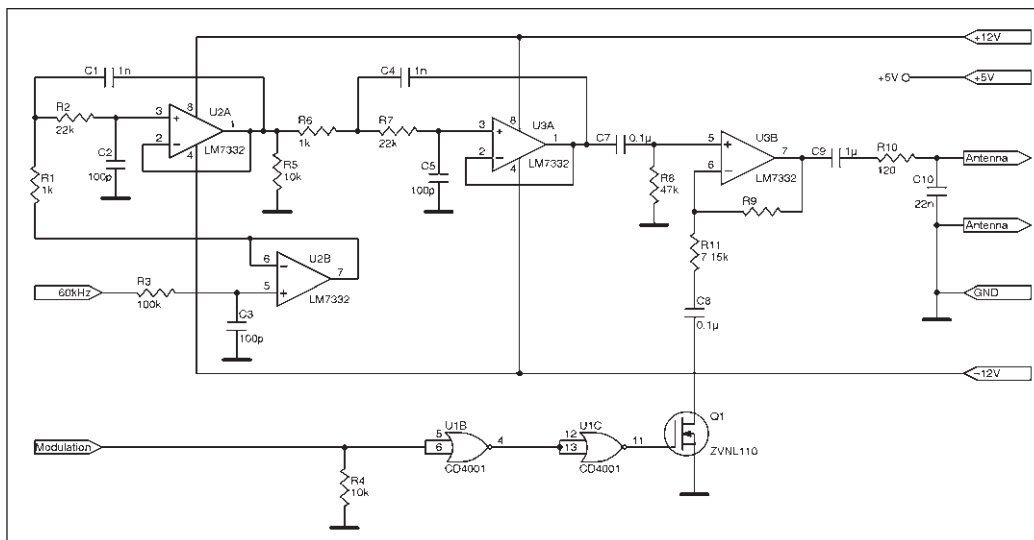
Texas Instruments, Inc. | [www.ti.com](http://www.ti.com)

### RESOURCES

J. P. Lowe and K. C. Allen, "Increasing the Modulation Depth of the WWVB Time Code to Improve the Performance of Radio-Controlled Clocks," The National Institute of Standards and Technology (NIST), 2006.

G. Novacek, "Sine Wave Synthesizer," *Circuit Cellar* 252, 2011.

—, "Filters, Filters Everywhere," *Circuit Cellar* 243, 2010.





**FIGURE 4**  
The WWVB transmitter modulator is built around a Texas Instruments LM7332 op-amp.

**PHOTO 2**  
The GPS clock standard includes an Adafruit Industries XBee transceiver.

GPS breakout board to build the GPS clock standard some time ago. It works very well right out of the box. My only complaint is that the SMA-to-uFL adapter, which connects the external antenna to the GPS module, is made of a stiff coaxial cable, which is incompatible with the tiny connector on the GPS board. Even though I provided a strain relief, the adapter eventually disconnected from the GPS. After reconnecting it several times, it ultimately destroyed the GPS connector and the entire module. I ordered a new module and a new adapter, connected them and, while they were lying free on the bench with no strain, I epoxied them together. It has been working well ever since. **Photo 2** shows the clock standard.

**NO TIME FOR DELAY**

I am happy to report that my “atomic” clocks no longer drive me crazy. However, I admit to one last shortcoming: there is an approximately 25-ms propagation delay between the GPS and the repeater. I considered compensating for it, but none of the WWVB clocks I have display seconds. Since I have already acclimatized myself to using a one-hand day-of-the-week clock, I decided to ignore that 25-ms delay. For now, at least. 



**ABOUT THE AUTHOR**

George Novacek 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. Contact him at gnovacek@nexicom.net with “Circuit Cellar” in the subject line.

COLUMNS


## ABOVE THE GROUND PLANE

# Strobe Photography

COLUMNS

Photographing a rapidly moving object requires careful synchronization of camera hardware, flash output, and object position. Ed explains how to sense the object, then trigger a camera and flash at the proper moments to capture interesting and useful images.

*By Ed Nisley (US)*



**PHOTO 1**

Flashing a white LED at 41.86 Hz freezes a fan turning at about 2,500 RPM. The camera integrated 20 flashes during the 0.5-s shutter time, with the aperture wide open at f/3.5, but the LED's low-light output still produced a dark image.

Although Hal "Papa Flash" Edgerton didn't invent the stroboscope, his high-speed flash inventions eventually produced sub-microsecond images of everything from fluid dynamics to nuclear explosions. The essential beauty of his iconic photographs came from synchronizing the camera and flash with an "interesting" event: a splashing milk drop, a shattering light bulb, or a vaporizing bomb casing.

In this article, I'll take a look at two computer-controlled photography applications: a strobe tachometer and a flash synchronized with a falling object. While my subjects aren't nearly as dramatic as Edgerton's, similar hardware and logic will let you capture useful images that, perhaps, have a touch of artistic merit.

## STROBE TACHOMETER

**Photo 1** shows a simple strobe tachometer in action: the white 10-mm LED in the upper left corner flashes once per fan revolution, about 2,500 revolutions and flashes per minute. Because each flash occurs when the blades are in the same position on each revolution, the camera shows a "stopped" fan, even though the breeze reminded me to keep my fingers clear. The camera shutter remained open for 0.5 s to integrate 20 flashes, which was also long enough for dim room light to expose the rest of the scene on my bench.

The pale square on the fan hub is retroreflective tape for my laser tachometer, which provided an independent measurement of the rotation speed. Because **Photo 1** shows a single square, you know that the flashes occur once per revolution; otherwise, you'd see multiple images of the tape. Although the LED would produce the same result if it flashed every second or third revolution, the picture would be even darker.

The schematic in **Figure 1** shows the MOSFET driver I used with an Arduino Pro Mini microcontroller to produce the pulses. The 10-mm LED contains five parallel white LED chips that require only 100 mA, so a small logic-level MOSFET without heatsinking suffices. Brighter LEDs drawing more current will also require a larger MOSFET and, depending on the pulse frequency and duty cycle, active cooling.

The Arduino Pro Mini hardware includes six PWM channels that produce fixed-frequency pulses with variable duty cycles, controlled by the `analogWrite()` function. However, a strobe tachometer requires a variable frequency and a fixed-pulse width, so the built-in function isn't useful. Instead, I set up hardware Timer1 to operate in Mode 14: Fast PWM, 16-bit resolution, pulse period set by ICR1, and pulse width set by OCR1A. **Listing 1** shows the code that initializes the timer hardware.

The timer includes a clock prescaler that divides the 16-MHz Arduino oscillator by several fixed ratios ranging from 1:1 to 1:1024. I picked the 1:64 ratio, so the timer clock ticks at 250 kHz = 16 MHz/64.

One tick of that clock would produce a 4- $\mu$ s LED pulse, which, while far longer than Doc Edgerton's strobe flash, emits so little light that neither my eyes nor my camera would notice it. Instead, I used a 100- $\mu$ s LED pulse that produces enough light to be visible. Converting 100  $\mu$ s into 25 timer ticks gives the value of `FlashLengthCt` in **Listing 1**.

Your application may require a shorter pulse or permit a longer one, depending on

the speed of the moving hardware. The tips of the 86-mm diameter fan blades in **Photo 1** traveled about 1 mm during each flash:

$$1.1 \text{ mm} = (\pi \times 86 \text{ mm}) \left(2,500 \frac{\text{rev}}{\text{min}}\right) \left(\frac{1 \text{ min}}{60 \text{ s}}\right) (100 \mu\text{s})$$

The fan blade tips appeared slightly blurred in real life, but you probably can't see that in the picture.


The maximum Timer1 period would be 64 K ticks = 262 ms = 3.8 Hz, with a nearly invisible 0.04% duty cycle. **Photo 1** required 41.7-Hz flashes at 0.4% duty cycle.

I used a rotary encoder knob to adjust the flash frequency in steps of either 1 or 0.01 Hz, selected by a push button switch in the knob shaft. The code in **Listing 2** updates the timer period based on the value of `KnobCounter`, a variable set by the interrupt handler to record the number of encoder clicks. The awkward calculations in the `printf()` statement that displays the pulse frequency work around the absence of floating-point format conversion, an understandable omission in a microcontroller language.

The program keeps the flash frequency between 4 and 2,500 Hz, which should synchronize with most vibrating and rotating objects. However, flash frequencies around 10 Hz may trigger vertigo or epileptic seizures in some people, so, if you're bothered by other flickering lights, you should modify the source code to exclude those frequencies.

Because there's no input from the rotating object, the LED flashes are not synchronized to the rotation, so you must adjust the flash frequency to match the rotation and stop its motion "by eye." Unlike the laser tachometer's digital output that I used for calibration, a strobe tachometer has a unique analog display: very small speed variations will be apparent as the object appears to rotate slowly in one direction or the other.

The complete source code available at *Circuit Cellar's* FTP site includes the interrupt



**FIGURE 1**

A small MOSFET suffices for a low-output LED strobe. Flashing high-output LEDs at higher duty cycles will require a more substantial driver.

```
analogWrite(PIN_STROBE,1); // let Arduino set up default Timer1 PWM
noInterrupts();
TCCR1B = 0; // turn off Timer1 for strobe setup
TCCR1A = 0x82; // clear OC1A on match, Fast PWM, lower WGM1x = 14
ICR1 = FlashPdCt; // set timer period
TCNT1 = FlashPdCt - 1; // force immediate update
OCR1A = FlashLengthCt; // set pulse duration
TCCR1B = 0x18 | TCCRxB_CS; // upper WGM1x = 14, Prescale 1:64, start Timer1
interrupts();
```

#### LISTING 1

These statements in the `setup()` function configure Timer1 to run in 16-bit fast PWM Mode 14, with ICR1 defining the period and OCR1A defining the pulse width. The CS bits select the 1:64 prescaler output from the Arduino Pro Mini's 16-MHz oscillator, so Timer1 ticks at 62.5  $\mu$ s.

```

if (KnobCounter) {
  noInterrupts();
  FlashFreq += (float)KnobCounter * FreqIncr;
  KnobCounter = 0;
  interrupts();

  FlashFreq = constrain(FlashFreq,FreqMin,FreqMax);
  FlashFreq = round(100.0 * FlashFreq) / 100.0;

  FlashPd = 1.0 / FlashFreq;
  FlashPdCt = FlashPd / TICKPD;

  noInterrupts();
  TCCR1B &= 0xf8;      // stop Timer1
  ICR1 = FlashPdCt;    // set new period
  TCNT1 = FlashPdCt - 1; // force immediate update
  TCCR1B |= TCCRxB_CS; // start Timer1
  interrupts();

  printf("Frequency: %d.%02d\n",
    (int)FlashFreq,
    (int)(100.0 * (FlashFreq - trunc(FlashFreq))));
}

```

**LISTING 2**

Turning the rotary encoder knob changes the flash frequency. When the `KnobCounter` variable becomes nonzero, this code updates the frequency, computes the corresponding period for the `Timer1` registers, and formats the new value to output to the Arduino serial monitor.

handler for the quadrature knob that controls the pulse frequency, but that's a simple matter of software.

**SYNCHRONIZED FLASHING**

Frustrated by the performance limitations of open-air and mercury arc lamps, Doc Edgerton invented a sealed flash tube filled with a noble gas, settling on Xenon because of its near-daylight emission spectrum. Today's miniature flash tubes descend directly from the multi-ton equipment he designed to photograph Normandy's beaches on June 5, 1944.

Even today, Xenon flash tubes produce more light per watt than high-output LEDs, by 20 or 30 dB. LEDs have more convenient form factors, use lower voltages, and integrate better with ICs, but they cannot compete when you need a short pulse of brilliant light. **Photo 2** captures a falling SD card case with a single 100- $\mu$ s flash from a modified Xenon flash unit, which shows an obvious improvement over **Photo 1**'s multiple flashes.

That photograph depends on several interacting components:

- Laser and photodiode beam-break detector
- Logic-level triggered Xenon flash unit
- Camera with motion detection
- LED to trigger the camera
- Arduino-based control

The flash must occur when the falling object is directly in front of the camera, which (for any given object) will be a known time after it interrupts the laser beam. The Arduino monitors the photodiode output and fires the flash at the appropriate moment, taking into account any flash trigger delay.


**Photo 3** shows a 3-D printed fixture that aligns a small red laser module with an IR LED used as a photodiode, with the laser focused near the center of the ring. Back in the 1960s, Forrest Mims discovered that LEDs can also act as narrow-bandwidth photodiodes with peak sensitivity to wavelengths somewhat shorter than they emit, which makes them useful sensors in ambient light conditions. In this case, the IR LED's bandwidth includes red laser light and excludes nearly all ambient light, so a simple heat shrink tubing shroud eliminates interference from the fluorescent fixture over my bench.

The laser module normally operates at 20 mA from a 3.3-V source, projecting a focused dot as far as the eye can see. The schematic in **Figure 2** shows that I reduced its current to 10 mA, just over the lasing threshold, with a 120- $\Omega$  series resistor from the 5-V supply and a PNP transistor as a power switch. The Arduino program turns the

**PHOTO 2**

Photographing a falling object in mid-air requires triggering a single, high-intensity flash at a precise moment. A modified Xenon camera flash produced a 100- $\mu$ s pulse 180 ms after this SD card passed through a laser beam detector.



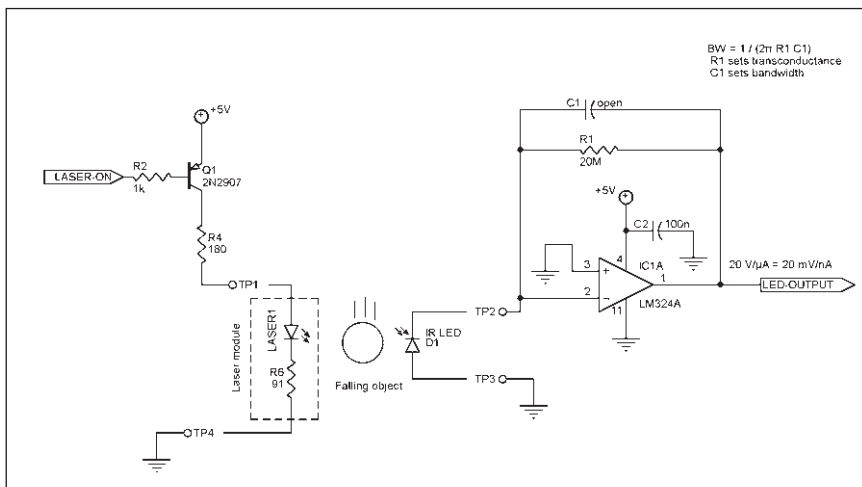


**PHOTO 3**  
This 3-D printed fixture aligns the laser module’s beam on an IR LED used as a photodiode. The beam is about 1-mm tall at the middle of the aperture. The laser beam’s backscatter produced the red glow on the fixture to the left of the brass laser module.

laser off after it detects an object, then turns it on after the timing sequence completes.


The LM324 op-amp in **Figure 2** acts as a transconductance amplifier to convert the photocurrent to a voltage for the Arduino’s input. Feedback resistor R1 provides enough gain that the reduced-strength laser beam saturates the LM324’s output at 4 V; R1’s actual value depends on the laser intensity, beam distance, and LED response, so the value you use will certainly differ from the 20 MΩ I used. A more modern single-supply and rail-to-rail op-amp would perform better, although this application has no need for high bandwidth or great accuracy.

I used a small relay to fire a surplus Xenon camera flash unit, driving the relay with another PNP transistor switch. The mechanical relay adds a 3-ms



**FIGURE 2**  
An external resistor reduces the laser module’s current to about 10 mA, just over its lasing threshold, to reduce its intensity. The IR LED operates as a photodiode sensitive to red light, with the op-amp converting the diode’s photocurrent to a voltage for the Arduino ADC input. The 20-MΩ resistor sets the transconductance ratio; its value depends on the laser intensity and photodiode sensitivity. Any single-supply op-amp should work in this circuit, with the LM324 definitely at the low-performance end of the scale.

# File System Solutions



COLUMNS

With **ALFAT SoC** You Can


**Use FAT/FAT32 with Long File Name**

**Use UART, SPI or I2C serial interface to access files.**

**Use SD-Reader Mode allowing the ALFAT SoC to act as an SD Card Reader.**


**Access files on two USB drives and an SD card simultaneously**

**Use high speed USB and 4-bit SD card interfaces.**



*The ALFAT SoC is also available as an OEM board.*

**GHIelectronics.com**



**PHOTO 4**

The control circuitry consists of an Arduino Pro Mini clone, three PNP transistors, an LM324 op-amp, a rotary encoder knob, and a few resistors. All the timing boils down to a simple matter of software!

delay between the Arduino's digital output and the light pulse, so the program must subtract that time from the overall delay. An optically isolated triac makes a better digital interface with zero delay, but the only optotriacs in my collection had too much leakage current and wouldn't fire the flash. The flash unit must have galvanic isolation from the rest of the circuitry, because its high voltages and pulsed currents pose hazards for both you and your hardware.

The flash unit always fires at full intensity, which produced enough light to illuminate a room and grossly overexpose objects a few hundred millimeters away. I replaced the 250- $\mu$ F electrolytic capacitor with a 1- $\mu$ F film capacitor, reducing the energy by nearly an order of magnitude, as the voltage increased from 300 to nearly 350 V. Remember that the energy stored in either of those capacitors can kill you.

**Photo 4** shows the breadboard layout, with a quadrature encoder knob on the far left and the Arduino Pro Mini with its USB interface near the middle, surrounded by the PNP transistors and the LM324 op-amp. The

Arduino Pro Mini runs from a 7.5-V power supply, but all the other circuitry uses a separate 5-V supply, because the Pro Mini's tiny SMD regulator cannot supply enough current at 5 V for anything else. The program sends timing and diagnostic information through the USB adapter, but you could add a small LED or LCD panel for that data. I don't have room here for the complete schematic, but you can download it and the Arduino program from *Circuit Cellar's* FTP site.

The quadrature knob sets the delay from the time the object breaks the laser beam until the flash fires, adding or subtracting 1 ms at each click. For most objects, you can get a reasonable time estimate from the ideal physics equation relating to the distance traveled under uniform acceleration to the elapsed time:

$$x = \frac{1}{2} at^2$$

$$t = \sqrt{\frac{2x}{a}}$$

I picked a convenient 200-mm drop distance that corresponds to 200 ms, after which I simply tuned the knob for the best picture.

It turned out that detecting the object and firing the strobe was the easiest part of the project. Triggering the camera shutter, which should have been trivially simple, posed the most problems.

## CAMERA SEQUENCING

The circuitry you've seen so far fires the flash after a fixed delay, presumably when the object reaches a specific position. The camera shutter must already be open when the flash fires, so the signal triggering the shutter must occur after the object breaks the laser beam and before the flash fires, with the exact time based on the delay from the signal to the shutter opening.

The Canon Hack Development Kit (CHDK) Free Software Foundation firmware for my Canon SX230HS camera adds a hardwired remote shutter release through the USB port, but some preliminary testing showed the delay was too unpredictable for my purposes. However, if you're taking ordinary pictures without tight timing requirements, a hardwired connection offers many advantages.

Somewhat to my surprise, the CHDK motion detection scripts had faster and more consistent timing after detecting a moving object, but they were also intended for ordinary photographic situations. My tests with small objects falling in front of a black background in a dim room didn't reliably



### ABOUT THE AUTHOR

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



# Join the Elektor Community

## Take Out a GOLD Membership Now!



### Your GOLD Membership Contains:

- 8 Regular editions of Elektor magazine in print and digital
- 2 Jumbo editions of Elektor magazine in print and digital (January/February and July/August double issues)
- Elektor annual DVD-ROM
- A minimum of 10% DISCOUNT on all products in Elektor.STORE
- Direct access to Elektor.LABS; our virtual, online laboratory
- Direct access to Elektor.MAGAZINE; our online archive for members
- Elektor.POST sent to your email account (incl. 25 extra projects per year)
- An Elektor Binder to store these projects
- Exclusive GOLD Membership card



### ALSO AVAILABLE:

The all-paperless GREEN Membership, which delivers all products and services, including Elektor magazine, online only.



Take Out Your Membership Now at [www.elektor.com/member](http://www.elektor.com/member)

Connect with us!



[www.facebook.com/elektorim](http://www.facebook.com/elektorim)



[www.twitter.com/elektor](http://www.twitter.com/elektor)

```

FallNow = millis() - FallStart; // assume falling: compute elapsed time

switch (Falling) {
case F_IDLE : // turn on laser for object detection
digitalWrite(PIN_LASER,LOW);
printf("Laser on, stabilizing... ");
while (analogRead(PIN_PHOTOCURRENT) <= DetectLevel) {
printf("**");
delay(1000); // ... throttle in case of blockage
}
printf("\nReady!\n");
Falling = F_WAIT;
break;
case F_WAIT : // record starting time of beam break
if (analogRead(PIN_PHOTOCURRENT) < DetectLevel) {
FallStart = millis();
Falling = F_DETECT;
}
break;
case F_DETECT : // turn off laser to signal detection
digitalWrite(PIN_LASER,HIGH);
Falling = F_PREFALL;
break;
case F_PREFALL : // turn on LED to trigger motion detection
if (FallNow >= DelayLED) {
digitalWrite(PIN_LED,LOW);
Falling = F_LED;
}
break;
case F_LED : // turn off LED
if (FallNow >= (DelayLED + PulseLED)) {
digitalWrite(PIN_LED,HIGH);
Falling = F_MD;
}
break;
case F_MD : // fire the strobe to take picture
if (FallNow >= DelayFlash) {
digitalWrite(PIN_FLASH,LOW);
Falling = F_FLASH;
}
break;
case F_FLASH : // turn off strobe relay
if (FallNow >= (DelayFlash + PulseFlash)) {
digitalWrite(PIN_FLASH,HIGH);
printf("Flash with LED delay: %d, Xenon delay: %d ...",
DelayLED,DelayFlash);
Falling = F_CLEAR;
}
break;
case F_CLEAR : // wait for camera to recycle
if (FallNow >= DelayClear) {
printf("done\n");
Falling = F_IDLE;
}
break;
default :
printf("** Bad Falling state: %02X",Falling);
Falling = F_IDLE;
}

```

**LISTING 3**

The Falling variable defines the state of the object as it passes through the hardware. After the laser sensor detects the object, all subsequent states occur after fixed durations, so there's no need for complex error handling.

trigger the script, so I added a white LED to the Arduino circuitry: flashing the LED reliably activated the motion detection script.

The LED appears in the lower right corner of **Photo 2** and the lower center of **Photo 5**. It's not illuminated, because the Arduino program turns the LED on and off before the shutter opens: a 50-ms pulse of light triggers the script, but 125 ms passes before the camera opens the shutter.


**Listing 3** shows the state machine that sequences all the operations required to capture an image of a falling object. After the initial detection in the `F_WAIT` state, each successive state change occurs after a fixed duration, so the state machine doesn't need elaborate controls and can't get locked up by a missed input. `F_CLEAR`, the final state, delays long enough to let the camera become ready for the next image, then returns to `F_WAIT` and turns the laser on again.


I combined eight separate exposures to produce **Photo 5** to show the effect of small variations in the release point. The Xenon flash on the right edge of the picture fired 200 ms after each ball broke the laser beam, but the balls appear in different positions because my fingertips didn't release them from exactly the same spot with the same (ideally zero) velocity. The lone ball above the cluster shows what happens with sticky clay: that one got off to a rather slow start after breaking the beam.

Sensing objects as they move past the sensor depends on a repeatable speed at that known position. I may change the setup so the Arduino program begins timing when it releases an object with an electromechanical actuator. The same type of state machine can control that output, as well as the camera and flash outputs, either in response to a Start button or upon sensing a Ready input.

Capturing a successful picture requires that everything occurs at exactly the right times and places, with precise repeatability and negligible jitter. That's obviously a mechanical problem, not a programming problem, so I foresee plenty of quality shop time in my future.

## CONTACT RELEASE

Hal Edgerton combined top-notch engineering with impressive photographic ability to produce spectacular images. This gadgetry will capture water-drop images and teach me a bit more about how he started it all. I can't imagine what he would have done with today's microcontrollers and digital cameras! 



**PHOTO 5**

This composite image of eight separate drops shows the effect of a few millimeters of fingertip inaccuracy and a nonzero release speed on the ball's position exactly 200 ms after the ball breaks the laser beam. The ball's speed is about 2 mm/ms.

## PROJECT FILES



[circuitcellar.com/ccmaterials](http://circuitcellar.com/ccmaterials)

## RESOURCES

The Edgerton Digital Collections, "Harold 'Doc' Edgerton," <http://edgerton-digital-collections.org>.

E. Nisley, "Pulsed LED Characterization," *Circuit Cellar* 278, 2013.

——, Softsolder.com, "Digitally Triggered Xenon

Flash," <http://softsolder.com>.

——, Softsolder.com, "Canon SX230HS vs CHDK: Motion-Detection Shutter Delay," <http://softsolder.com>.

Wikia, "What is CHDK?" <http://chdk.wikia.com>.

Wikipedia, "Flicker vertigo."

——, "Forrest Mims."

## SOURCE

**Arduino Pro Mini**

Arduino | <http://arduino.cc>

## FROM THE BENCH

# LCD User Interface

## Reconfiguring and Serial Backpacking

COLUMNS



By Jeff Bachiochi (US)

Many personal and consumer projects rely on small-character LCDs. Adding serial backpacking can create a more user-friendly interface. This article discusses some options for enhancing LCDs.

**B**efore microcontrollers, all peripherals were connected to a system's external address space and 8-bit (and larger) buses were the norm. Whether your system had separate address and data buses or used a multiplexed bus, parallel traces ran all around PCBs.

Obviously, these are still being used in higher-power computing devices (e.g., cell phones, tablets, and PCs). For less-demanding applications, we can enjoy the simplicity that the microcontroller brings.

Moving data in byte (word or larger) chunks gets the job done quickly (if there are equal execution times) when compared to moving it serially. Throughput is high when you can simultaneously transfer all the data. However, many microcontroller applications do not require such throughput, so we must bow to these older technologies to use some peripherals. The LCD is one of those peripherals stuck in the last century. It seems manufacturers feel it is one of those "if it ain't broke don't fix it" situations.

LCDs usually have a minimum of 14 connections including an 8-bit data bus with several control signals in addition to power and ground. Let's start with a review of how to communicate with these devices. The 8-bit bus enables a full set of ASCII characters to be selected. Control characters can be replaced

with user-defined glyphs (any combination of the typical 5 × 7 dot blocks), which gives the user some artistic control. A display has two separate areas of interest. There are obviously characters to display (data area) and there is some control over where and how these are displayed (control area). The register select (RS) signal tells the display which of these areas will use the 8-bit data.

Once initialized, you just write data to display. However, the display can provide some feedback to the user (i.e., when it is too busy to accept new data), so a read/write (R/WR) signal is also provided. Finally, once these signal states and the data have been selected, the E (control strobe) signal is used to indicate that the display should perform the requested task. **Figure 1** shows the signal timing required to send or receive one byte.

Each display uses a processor to handle this interface and control the LCD segments in a multiplexed arrangement. Most displays use the same processor, regardless of the display format, from one line of eight characters to four lines of 32 characters. The display character positions are pre-mapped into 128 bytes of memory, 64 bytes for each of two lines. In a 16 × 2 display, the second row begins at address 0x40 and not 0x10 (the next logical address). You can write characters to addresses 0x10-0x3F without

a problem, but you will only see these if the two-line LCD supports a 0x40-character line length. If the display has four lines, then the third and fourth lines are actually extensions of the first and second lines. So, at the end of the first display line (0x1F) the third display line begins (0x20). It's a bit convoluted, yet logical if you think of how the display is pre-mapped.

Once the display exceeds 128 characters, a second E signal is used, as there are two processors, each with its own spaces. There are 256 different instructions that can be written to the processor's command port (see **Table 1**). Half of these (1xxxxxx) are taken up by the set DDR RAM Address command. Half of the remainder (01xxxxxx) are taken up for the set CG Address command.

The rest of the commands (00xxxxxx) are used for various other functions. One of these functions enables an auto increment of the address pointer after each write to the data space. This eliminates the need to alternately write to the control space to set the address and write to the data space to display a character. Its configuration also enables moving the display window, which looks like the text is scrolling across the display.

There are also functions such as the Clear Display command. A write to the command port with the 8-bit value 0x01 will clear the display. These commands help customize the display's operation. They must be appropriately set for the display processor to properly handle the display. They are normally set once on power-up to put the processor on the right path. This initialization sequence must follow a particular timing regimen until the processor has enough information to intelligently reply with a busy/not busy status. **Table 1** shows the instruction set for a typical display. Note the state of each control signal as well as the 8-bit value are shown.

### SQUASHING THE 11-BIT INTERFACING

When using an LCD with a microcontroller, I/O pins are always at a premium. The microcontroller you choose is often based on the minimum number of I/Os required for an application. Because microcontrollers have many peripherals built in, the need for wide external buses disappeared. The requisite 36-plus I/O pins went away and the size of microcontrollers shrank when smaller amounts of I/O were needed. When some manufacturers came out with microcontrollers with eight (and even six) pins, there was much controversy. But that's another story.

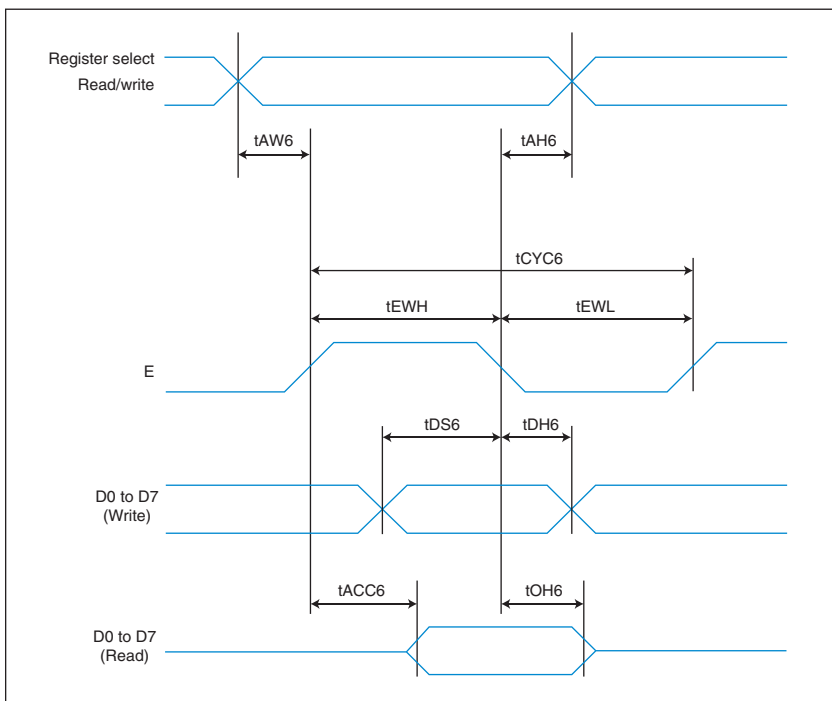
Having to set aside 11 pins just to use an LCD was a constant concern. Most LCD manufacturers have an alternative to the 8-bit

data bus. Without redefining the interface, you can configure the device to use a 4-bit bus. This comes with a compromise in speed because each transaction then requires two nibble transfers. The two transfers must be in most-significant nibble/least-significant nibble (MSN/LSN) format.

The most perplexing issue for many is the initialization of 4-bit mode. If there are only four data lines connected, how can this work? The use of 4-bit mode requires connection to the high bits (bit 7:4). A write in 4-bit mode will be a 0x0-0xF value. This will look like "0x0?-0xF?" because you have no control over the lower bits.

If you refer back to **Table 1**, you will see that the Function Set command uses bit 4 to configure the bus width. An 8-bit value of 0x2? will configure the display for a 4-bit bus. An 8-bit value of 0x3? will configure the display for an 8-bit bus. Note that since you don't have control of the lower bits, the other functions in this command could arbitrarily be set to anything. The problem is not knowing the display processor's present mode. Is it in 8-bit mode or 4-bit mode? Is it expecting a byte, an upper nibble, or a lower nibble?

To take control of the situation, you must send the 0x3 value multiple times. If the processor is expecting a byte, it will receive 0x3? (and it will be content with each transfer). Remember, you are only sending



**FIGURE 1**

When read/write (R/WR) = 1, the register select (RS) control states must be set up before the E (control strobe) signal is used to request 8-bit data. Likewise, when R/WR = 0, the RS and the 8-bit data must be set up before the E (control strobe) signal is used to signal the processor it can capture the data.

Instruction	Code										Description	Execution Time (Maximum) When $F_{CP}$ or $F_{OSC}$ Is 250 kHz	
	RS	R/W	DB <sub>7</sub>	DB <sub>6</sub>	DB <sub>5</sub>	DB <sub>4</sub>	DB <sub>3</sub>	DB <sub>2</sub>	DB <sub>1</sub>	DB <sub>0</sub>			
Clear display	0	0	0	0	0	0	0	0	0	0	1	Clears entire display and sets DDR RAM address 0 in address counter.	15.2 ms
Return home	0	0	0	0	0	0	0	0	0	1	x	Sets DDR RAM address 0 in address counter. Also returns shifted display to original position. DDR RAM contents remain unchanged.	15.2 ms
Entry mode set	0	0	0	0	0	0	0	0	1	I/D	S	Sets cursor move direction and specifies shift or display. These operations are performed during data write and read.	40 $\mu$ s
Display on/off control	0	0	0	0	0	0	0	1	D	C	B	Sets on/off of entire display (D), cursor on/off (C), and blink of cursor position character (B).	40 $\mu$ s
Cursor or display shift	0	0	0	0	0	0	1	S/C	R/L	x	x	Moves cursor and shifts display without changing DDR RAM contents.	40 $\mu$ s
Function set	0	0	0	0	0	1	DL	N	F	x	x	Sets interface data length (DL), number of display lines (N), and character font (F).	40 $\mu$ s
Set CG RAM address	0	0	0	0	1	ACG					Sets CG RAM address. CG RAM data is sent and received after this setting.	40 $\mu$ s	
Set DDR RAM address	0	0	1	ADD							Sets DDR RAM address. DDR RAM data is sent and received after this setting.	40 $\mu$ s	
Read busy flag (BF) and address	0	1	BF	AC							Reads busy flag (BF) indicating internal operation is being performed and reads address counter contents.	40 $\mu$ s	
Write data to CG RAM or DDR RAM	1	0	Write data								Writes data into DDR RAM or CG RAM.	40 $\mu$ s	
Read data from CG RAM or DDR RAM	1	1	Read data								Reads data from DDR RAM or CG RAM.	40 $\mu$ s	
	I/D = 1: Increment I/D = 0: Decrement S = 1: Accompanies display shift S/C = 1: Display shift S/C = 0: Cursor move R/L = 1: Shift to the right										DDR RAM: Display data RAM CG RAM: Character generator RAM ACG: CG RAM address ADD: DDR RAM address Corresponds to cursor address AC: Address counter used for both	Execution time changes when frequency changes (e.g., when $F_{CP}$ or $F_{OSC}$ is 270 kHz: 40)	
	R/L = 0: Shifts to the left DL = 1: 8 bits, DL = 0 : 4 bits N = 1: 2 lines, N = 0 : 1 line F = 1: 5 $\times$ 10 dots, F = 0: 5 $\times$ 7 dots BF = 1: Internally operating BF = 0: Can accept instruction										DDR RAM and CG RAM address	40 ms $\times$ 250/270 = 37 ms	
x = don't care (no effect)													

TABLE 1

This is the command table for an Optrex parallel LCD. Note the commands follow a 2N relationship (0b00000001, 0b0000001x, 0b000001xx..., etc.). When RS = 0, the first half (0x00-0x7F) are control functions and the second half (0x80-0xFF) is the address pointer to the display memory addresses. When RS = 1, the value gets stored at the address pointer's display address or is read from the address.

**Function Set Command**

	Bit 7	Bit 6	Bit 5	Bit 4	Bit 3	Bit 2	Bit 1	Bit 0
Standard register set	0	0	1	DL	N	F	0	0
Use alternate set 1	0	0	1	DL	N	F	0	1
Use alternate set 2	0	0	1	DL	N	F	1	0
Not used at this time	0	0	1	DL	N	F	1	1

**TABLE 2**

Bits 1 and 0 of the Function Set command define which register set is being accessed.

4-bit data to the 8-bit LCD data bus' upper nibble. If it is expecting an upper nibble, it will receive the 0x3 and hold on to this as the MSN and look for a lower nibble. The next 0x3 will be received as the LSN and be satisfied with the command 0x33 (use 8-bit mode next time).

A subsequent transfer will receive 0x3? (and it will be happy since it is in byte mode). If it is expecting a lower nibble, the 0x3 will be received as the LSN. However, you can't be sure what it has for an MSN! At this point, the only thing you can be sure of is that the next transfer *isn't* a least-significant bit (LSB). Therefore, one of the previous rules will be satisfied on subsequent transfers. So, the worst-case scenario is that you will need to send 0x3 three times to ensure you end up in

8-bit mode. During this process, you must use a timing loop to ensure the display processor is ready for the next transfer, as the busy bit can't be accessed yet.

Once the display processor has been forced into 8-bit mode by sending individual transfers of 0x3, you can send one 0x2 transfer (remember it's in 8-bit mode now). This is received as a complete 0x2? command. The processor will then switch to 4-bit mode and will expect two nibble transfers for each instruction. Since you had no control over the lower 4 bits of that last (8-bit) command, you must again write to the function set. This time, you will have control of the lower nibble value to configure these lower bits as necessary (refer back to **Table 1**).

Obviously transfers will be quicker (and

**( Ta = -35°C to 85°C )**

Item	Signal	Symbol	Condition	V <sub>DD</sub> = 2.7-to-4.5-V rating		V <sub>DD</sub> = 4.5-to-5.5-V rating		Units
				Minimum	Maximum	Minimum	Maximum	
Serial clock period	SCL	t <sub>SCYC</sub>	-	200	-	100	-	ns
SCL "H" pulse width		t <sub>SHW</sub>		20	-	20	-	
SCL "L" pulse width		t <sub>SLW</sub>		160	-	120	-	
Address setup time	RS	t <sub>SAS</sub>	-	10	-	10	-	ns
Address hold time		t <sub>SAH</sub>		250	-	150	-	
Data setup time	SI	t <sub>SDS</sub>	-	10	-	10	-	ns
Data hold time		t <sub>SDH</sub>		10	-	20	-	
CS-SCL time	CS	t <sub>CSS</sub>	-	20	-	20	-	ns
		t <sub>CSH</sub>		350	-	200	-	

\*1 All timing is specified using 20% and 80% of V<sub>DD</sub> as the standard.

**TABLE 3**

The SPI is quite fast. Even with a 6' extension cable, I had no trouble conversing with the LCD controller/driver.

Function set	0	0	0	0	1	DL	N	DH	IS2	IS1	DL: interface data is 8/4 bits. N: number of line is 2/1. DH: double height font IS[2:1]: instruction table select	26.3 $\mu$ s
--------------	---	---	---	---	---	----	---	----	-----	-----	---	--------------

### Instruction table 0 (IS[2:1] = [0,0])

Cursor or display shift	0	0	0	0	0	1	S/C	R/L	X	X	S/C and R/L: set cursor moving and display shift control bit, and the direction, without changing DDR RAM data.	26.3 $\mu$ s
Set CG RAM	0	0	0	1	AC5	AC4	AC3	AC2	AC1	AC0	Set CG RAM address in address counter.	26.3 $\mu$ s

### Instruction table 1 (IS[2:1] = [0,1])

Bias set	0	0	0	0	0	1	BS	1	0	FX	BS = 1:1/4 bias BS = 0:1/5 bias FX: fixed on high in three-line application and fixed on low in other applications	26.3 $\mu$ s
Set ICON address	0	0	0	1	0	0	AC3	AC2	AC1	AC0	Set ICON address in address counter.	26.3 $\mu$ s
Power/ICON control/contrast set	0	0	0	1	0	1	Ion	Bon	C5	C4	Ion: ICON display on/off Bon: set booster circuit on/off. C5,C4: contrast set for internal follower mode.	26.3 $\mu$ s
Follower control	0	0	0	1	1	0	Fon	Rab 2	Rab 1	Rab 0	Fon: set follower circuit on/off. Rab2~0: select follower amplified ratio.	26.3 $\mu$ s
Contrast set	0	0	0	1	1	1	C3	C2	C1	C0	Contrast set for internal follower mode.	26.3 $\mu$ s

### Instruction table 2 (IS[2:1] = [1,0])

Double-height position select	0	0	0	0	0	1	UD	X	x	x	UD: Double height position select	26.3 $\mu$ s
Reserved	0	0	0	1	X	X	X	X	X	X	Do not use (reserved for test).	26.3 $\mu$ s

**TABLE 4**

Here is the shortened version of the command table for an Electronic Assembly DOG LCD. I've included the function set as well as the extended registers. The extended or alternate registers are selected using bits 1:0 of the function set.



easier) once you can depend on the use of the busy status bit instead of a fixed-minimum delay time. Up to this point you've only written to the command port. Reading from the command port will return the display's status. Bit 8 of this value is the busy bit. A "1" indicates the processor is busy with housekeeping chores and you should refrain from sending it any data. The other seven bits are the address pointer's present value, which is 0x00-0x7F. This value represents the character position 0-127. Remember to stay in sync. In 4-bit mode you must do two reads and combine the nibbles to get a byte value.

This reduces the interface from 11 bits to seven bits and is a good compromise. However, many feel the power of a microcontroller is in its serial peripherals. Can these be used to reduce an LCD interface even further?

### BACKPACKING

Some of you may be familiar with the term "serial backpack." This refers to a PCB interface that goes between a parallel LCD and your microcontroller. It incorporates an additional microcontroller that directly connects to the display in 8-bit (or 4-bit) mode and provides a serial I/O connection as



the user interface. These can cost as much as the display itself. While they are a great fix, it's time for manufacturers to play catch-up and bring our little friends into the 21st century.

Electronic Assembly's DOG family of LCDs are high-contrast, easy-to-read, super-twist

#### PHOTO 1

Electronic Assembly's LCD Tools for Windows demonstration mimics a real display on your PC screen. You can change display type and backlight colors. I've displayed a sign-on message that is used with this project.

COLUMNS



Sign up today and **SAVE 50%** • Sign up today and **SAVE 50%** • Sign up today and **SAVE 50%**

## Now offering student SUBSCRIPTIONS!

When textbooks just aren't enough, supplement your study supplies with a subscription to *Circuit Cellar*. From programming to soldering, robotics to Internet and connectivity, *Circuit Cellar* delivers the critical analysis you require to thrive and excel in your electronics engineering courses.

Sign up today and **SAVE 50%**

Sign up today and **SAVE 50%**

Sign up today and **SAVE 50%**

[www.circuitcellar.com/subscription](http://www.circuitcellar.com/subscription)



**PHOTO 2**

The Electronic Assembly DOG display I prototyped fits into a small plastic case. The project's prototype along with a chip on glass (COG) display and backlighting module are shown. Note at least the four lower pins (unused by the display) must be soldered to the backlight module to bring backlight power to the module.

display modules that require 3.3 or 5 V to operate and include a SPI-type interface. I learned about these while looking for a small-character display that had red backlighting. I wanted something that could be used at night without inhibiting night vision. Night vision happens when a lack of light entering our eyes signals the iris to open up and let in more light. However, even the shortest exposure to a light source will close the iris, hindering the ability to "see well" in low-light conditions. It seems that exposure to red light is the least disruptive and won't completely ruin night vision. If you look into Electronic Assembly's complete line of character and graphic displays, you'll find support for many serial interfaces as a standard offering.

The DOG family is very versatile. The chip on glass (COG) displays are completely separate from any backlight. They are available in both character and graphic models with transmissive (STN) or transreflective (FSTN) technologies in either positive or negative formats. Backlighting is provided through a piggyback module that goes between the display and your PCB. These are available in red, white, blue, green, yellow, amber, and red/green/blue (RGB). Note: The DOG graphic displays are monochrome just like the character displays. For full-color displays look at the DIP family.

**FOUR-WIRE SPI**

The SPI consists of a synchronous clock (SCL) and two data lines, one in and one out. Normally labeled as SDI and SDO, these data paths are often confused because they are defined as the direction of the data by the host, but devices sometimes use the same labels. I prefer using the MOSI and MISO nomenclature, as it more fully describes the signal connection. Many SPI peripherals also include a select line (chip select). You can use chip select to enable a single device from multiple devices that all share the clock and data lines.

Electronic Assembly has chosen to substitute an RS input for the MISO signal. This means two things. First, there is no feedback or reading data from the device (negative consequence: no busy bit access). Second, since RS has a direct connection, no extra data has to be sent to handle this bit (positive consequence: throughput is doubled).

Sitronix Technology's ST7036 LCD controller/driver's datasheet shows a complete description of the available formats, which include I<sup>2</sup>C. This format requires the RS state to be passed as an extra command byte. While the I<sup>2</sup>C format is not available on this particular LCD, it suggests one of the compromises necessary when reducing the required I/Os.

This modified SPI does not require any extraneous data to be passed. The same data values are passed that would be placed on an 8-bit parallel interface. In fact, there are a few additions to the standard command table. I'll get to these in just a bit. First, why would any new commands be needed?

These displays have built-in bias circuitry and are contrast adjusted. The displays can be run at 5 or 3.3 V. The bias requires changes depending on which supply voltage is used. So, a few additional registers are implemented to handle making adjustments to this circuitry. Referring back to **Table 1**, you may notice a few "x" (don't care) register bits. Specifically, in the Function Set register, bits 1:0 can be any state. With this new LCD, these bits take on meaning (see **Table 2**).

To use alternate register sets, the user must first send a Function Set command with the two LSBs set to the chosen register set. Accessing alternate register sets is normally only necessary during initialization. Once any alternate registers have been correctly set, the user should select the standard register set for normal use.

One feature of using this (multi-line) display is double-height characters, which includes combining character heights for



circuitcellar.com/ccmaterials

**SOURCES****DOG LCD Family and LCD Tools for Windows**

Electronic Assembly GmbH | [www.lcd-module.de](http://www.lcd-module.de)

**ST7036 LCD Controller/driver**

Sitronix Technology Co., Ltd. | [www.sitronix.com.tw](http://www.sitronix.com.tw)

the three-line display. When I viewed the controller datasheet, I got excited when I saw the double-height characters drawn without a break (i.e., as contiguous pixels, when most LCDs have a row of non-pixels separating lines of which you have no control). When I finally connected the display, I found the double-height characters do in fact have a blank row of pixels between the characters top and bottom halves.

Referring to **Table 3**, you can see that at 5 V the DOG display can accept data at a more than one data byte per microsecond rate. While this throughput is fast, you still need to pay attention to how quickly the display can internally execute a particular command. When using a display in the 8- and 4-bit parallel mode, you can perform writes and reads on the display. In these modes you can read the display processor's status to determine when it is ready for new data.

In fact, even when using a parallel interface, upon power-up, you must place delays between data bytes to give the display a chance to get itself initialized to the point where it can respond with status information (as mentioned earlier). This same delay procedure must be used between all bytes when using the SPI. The command table provides times for the execution of each command. For the DOG displays, this time is 26.3  $\mu$ s for every command except power up, which is power-supply dependent, and the Clear Display and Return Home command, which is 1.08 ms (see **Table 4**).

I set up a timer to deliver a 26.3- $\mu$ s timeout. My delay routine uses a `Loops` byte to determine how many timer overflows to count before returning. When called with `Loops=1`, a single timeout is used. By changing `Loops` I can get multiples of 26.3  $\mu$ s. For example `Loops=41` produces 1.078 ms (i.e., 26.3  $\mu$ s  $\times$  41). This keeps delays to an absolute minimum and dependent on the command being issued.

## BACKLIGHT

Electronic Assembly's LCD Tools for Windows is a free program that shows the different displays and backlighting colors you can choose (see **Photo 1**). Most of today's displays are purchased with backlighting as an option. This means you are typically stuck with whatever is available for that model. The DOG backlight is purchased separately. This provides more (color) options than with any other display. An RGB option even gives you control over the color. The display glass has 20 pins for the display interface along the top, but also a few pins on its lower edge that provide a rigid mounting to your PCB (see **Photo 2**).

Backlighting modules are assembled on

0.031"-thick PCBs that have feedthrough holes for all of the display's pins. The backlight module uses the display's lower (unused) pins to receive power from your PCB's circuitry. These pins require a solder connection on the LED module. You must add current limiting (resistors) to the backlighting LEDs depending on the power supply voltage (5/3.3 V). These two circuits must be paralleled when using 3.3 V but can be connected in series when using 5 V. Note that the RGB module gives you control over each of the three LED colors.


I like being able to choose between positive (black characters on a colored background) or negative (colored characters on a black background) displays. Because of COG technology, these displays are physically smaller than most parallel displays.

I mounted my display in a small plastic Hammond case (see **Photo 2**). With a six-pin connector on the end of the case, I can mount this display at some length from its controller.

The same flexibility provided in this one-, two-, and three-line display series is given to larger character displays as well as the full-color graphic displays with touch overlays offered by Electronic Assembly. I wish there was a more complete description of the SPI.

The ST7036's datasheet provides a lot of information explaining the I<sup>2</sup>C interface and little information about the SPI. I would like to see more about the modified SPI using the RS line.

Might this be the beginning of a new era in low-cost LCDs? Small-character displays are the mainstay of many personal and consumer projects. It's about time that they are given a makeover. Manufacturers should pay attention to how engineers have added backpacking to give these displays a more user-friendly interface. I'm talking about driver manufacturers here.

Unless we go back to using controller/drivers—where this may be more easily implemented—we are adding unnecessary costs by having to include an intermediate interface translator (i.e., microcontroller). Fortunately, this trend is beginning to take hold, and is available today from a few manufacturers. 

## ABOUT THE AUTHOR


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](mailto:jeff.bachiochi@imaginethatnow.com) or at [www.imaginethatnow.com](http://www.imaginethatnow.com).



# CROSSWORD

**JULY 2014**

The answers will be available at [circuitcellar.com/crossword](http://circuitcellar.com/crossword).



## ACROSS

2. Inactive but still available
4. Used for circuitry interfacing [two words]
7. Open-source web tool developed by members of the World Wide Web Consortium (W3C)
8. A continuous sequence [two words]
9. A tower PC, for example [two words]
11. Explosive code [two words]
15. A software glitch that changes its conduct when analyzed
16. Makes things appear to move slowly or not at all
17. Approximately 0.333 NA
18. Unix command-line interpreter developed by and named after a Bell Labs employee [two words]
19. Defines a point in 3-D

## DOWN

1. Signal processing for sensor arrays
3. Works in double-duplex mode [two words]
4. Unix-based command-line utility
5. Used to convert to intermediate frequencies
6. Creates data words
10. Uses solar power to create energy
12. File sharing protocol
13. E.g., gibi [two words]
14. Linux distribution based on Slackware

**What's your EQ?** The answers are posted at [www.circuitcellar.com/category/test-your-eq/](http://www.circuitcellar.com/category/test-your-eq/). You can contact the quizmasters at [eq@circuitcellar.com](mailto:eq@circuitcellar.com).

# TEST YOUR EQ

*Contributed by David Tweed*

## PROBLEM 1

When designing a pair of band-splitting filters (for, say, an audio crossover), why is it important to match frequencies of the  $-3$ -dB points of the low-pass and high-pass responses?

## PROBLEM 2

A certain portable stereo unit runs for about 12 h on a set of LR20 (D-size alkaline) batteries. If you want to extend the stereo's runtime, is it better to simply use multiple sets of batteries sequentially, or to connect them all in series-parallel to create one big battery pack?

## PROBLEM 3

If you wanted to make a capacitor from scratch, what common household materials might you use?

## PROBLEM 4


How big would a  $10$ - $\mu$ F capacitor using these materials be?

engineering embedded  
 DISCUSS social media  
 information design tips tutorial software  
 engineering tools  
 audio business talk  
 electronics media  
 system data COMMUNITY  
 NETWORKING  
 product news  
 projects engage

Want to talk to us directly?  
 Share your interests and opinions!  
 Check out our New Social Media Outlets  
 for direct engagement!

**CIRCUIT CELLAR / AUDIOXPRESS / ELEKTOR**

## CC SHOP

**1 PC PROGRAMMING**

Many different PC programming languages are available. Some have beautiful names, some have easy-to-use development tools, others have incredible power. They all have one thing in common: they assume that you have, or want to have, a knack for technology and difficult-to-read commands. Using SmallBASIC, you will have an application up and running in a matter of minutes.

*Author: Bert van Dam*

*Item #: BK-ELNL-978-1-907920-26-4*

**4 RFID MIFARE AND CONTACTLESS CARDS IN APPLICATION**

RFID technology is now being used in many areas in which barcodes, magnetic strips, and contact smart-cards were previously used. This book provides a practical and comprehensive introduction to MIFARE, which is the most widely used RFID technology. The initial chapters cover physical fundamentals, relevant standards, RFID antenna design, security considerations, and cryptography.

*Author: Gerhard H. Schalk and Renke Bienert*

*Item #: BK-ELNL-978-1-907920-14-1*

**2 ANDROID APPS: PROGRAMMING STEP-BY-STEP**

Many smartphones and tablet computers are powered by an Android OS. These portable devices' speed and computing power enable them to run applications that would have previously required a desktop PC or custom-designed hardware. *Android Apps* introduces you to the programming required to design apps for Android devices. Operating the Android system is explained step-by-step to show how personal applications can be easily programmed.

*Author: Stefan Schwark*


*Item #: BK-ELNL-978-1-907920-15-8*

**3 CC VAULT**

CC Vault is a pocket-sized USB that comes fully loaded with every issue of *Circuit Cellar* magazine. This comprehensive archive provides an unparalleled amount of embedded hardware and software design tips, schematics, and source code. CC Vault contains all the trade secrets you need to become a better, more educated electronics engineer.

*Item #: CCVAULT*

**CC SHOP**



**5 LEARNING TO FLY WITH EAGLE V6**


EAGLE is a user-friendly, powerful, and affordable software package for efficient PCB design. It can be used on most main computing platforms including: Microsoft Windows, Linux, and Apple Mac OS X. *EAGLE V6* will benefit novices and professionals who are eager to learn about EAGLE or may be migrating from another computer-aided design (CAD) package. From schematic and layout editing tools to project completion, this book will help you achieve your PCB fabrication goals.

*Author: Mitchell Duncan*  
*Item #: BK-ELNL-978-1-907920-20-2*

**8 MASTERING MICROCONTROLLERS: HELPED BY ARDUINO**

Arduino boards have become hugely successful. They are simple to use and inexpensive. *Mastering Microcontrollers* will teach you how to program microcontrollers and help you turn theory into practice using an Arduino programming environment. Become a master today!


*Author: Clemens Valens*  
*Item #: BK-ELNL-978-1-907920-23-33*



**6 PRACTICAL DIGITAL SIGNAL PROCESSING USING MICROCONTROLLERS**

Digital signal processing (DSP) reflects the growing importance of discrete time signals and their use in everyday microcontroller-based systems. In this book, the author presents the basic theory of DSP with minimum mathematical treatment and teaches you how to design and implement DSP algorithms using popular PIC microcontrollers. You'll be up and running in no time!

*Author: Dogan Ibrahim*  
*Item #: BK-ELNL-978-1-907920-21-9*



**7 CIRCUIT CELLAR ISSUE PDFS**

Have you ever missed an issue of *Circuit Cellar* and didn't know where to find it? Grab a digital version! From 1998 to the present, dive into more than 100 issues of embedded electronics insights that you may have missed. Can you envision the evolution and development across a decade? With technology advancing every day, don't fall behind. Learn about wearable wireless transceivers from Mathew Laibowitz and Joseph Paradiso or a DIY single-board computer (SBC) from Oscar Vermeulen and Andrew Lynch. The topics and projects are endless.

Further information and ordering

[www.cc-webshop.com](http://www.cc-webshop.com)  
**CONTACT US:**

Circuit Cellar, Inc.  
 111 Founders Plaza, Suite 300  
 East Hartford, CT 06108  
 USA  
 Phone: 860.289.0800  
 Fax: 860.461.0450  
 E-mail: [custservice@circuitcellar.com](mailto:custservice@circuitcellar.com)

# IDEA BOX

## the directory of PRODUCTS & SERVICES

### AD FORMAT:

Advertisers must furnish digital files that meet our specifications (circuitcellar.com/mediakit).

ALL TEXT AND OTHER ELEMENTS MUST FIT WITHIN A 2" x 3" FORMAT. E-mail [adcop@circuitcellar.com](mailto:adcop@circuitcellar.com) with your file.

For current rates, deadlines, and more information contact Peter Wostrel at 978.281.7708 or [circuitcellar@smmarketing.us](mailto:circuitcellar@smmarketing.us).

The Vendor Directory at [circuitcellar.com/vendor](http://circuitcellar.com/vendor) is your guide to a variety of engineering products and services.

**"Your reliable outsourcing partner"**

**LOW cost**  
with **HIGH quality**  
**PCB, PCBA and More**

INSTANT QUOTE AT:  
[www.myropcb.com](http://www.myropcb.com)

OR CALL:  
**1-888-PCB-MYRO**



**Internet Framework  
Just a Click Away**

**PIC® MCU WIZARDS**

**CCS C Compiler!**

Easily Start projects for  
PIC® Peripherals including  
USB & TCP/IP...and more

Add a Internet Web Server  
& Send Emails with just  
a few clicks!

sales@ccsinfo.com  
262-522-6500 x 35



**GHz BGA/QFN Sockets  
0.3mm to 1.27mm**

**Industry's Smallest Footprint**

- Up to 500,000 Insertions
- Bandwidth up to 40 GHz
- 2.5mm per side larger than IC
- Ball Count over 3500
- Body Size 2 - 100mm
- Five different contactor options
- Optional heatsinking to 100W
- Six different Lid Options
- <25 mΩ Contact Resistance throughout life




RoHS

Ironwood  
ELECTRONICS 1-800-404-0204  
[www.ironwoodelectronics.com](http://www.ironwoodelectronics.com)

**ALL ELECTRONICS CORPORATION**

**Electronic and Electro-mechanical Devices, Parts and Supplies.**  
Many unique items.

We have what you need for your next project.



[www.allelectronics.com](http://www.allelectronics.com)  
Free 96 page catalog 1-800-826-5432

**ADuC841  
Microcontroller  
Design Manual**

From Microcontroller Theory to Design Projects

Shop for this book, and others, at  
[www.cc-webshop.com](http://www.cc-webshop.com)



**FPGA Boards from JAPAN** **HUMAN DATA**

**SAVING COST=TIME with readily available FPGA boards**

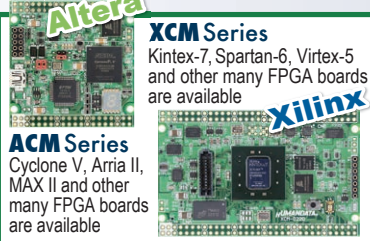
**Altera**  
**XCM Series**  
Kintex-7, Spartan-6, Virtex-5 and other many FPGA boards are available

**Xilinx**  
**ACM Series**  
Cyclone V, Arria II, MAX II and other many FPGA boards are available

- Basic and simple features, single power operation
- Over 100 varieties of FPGA/CPLD boards are available

**PLCC68 Series**

- Designed for 68-pin PLCC socket
- Very small size (25.3 x 25.3 [mm])
- 50 I/Os (External clock inputs available)
- 3.3V single power supply operation



[www.hdl.co.jp/CC/](http://www.hdl.co.jp/CC/)



Join The  
**INTERNET of THINGS**  
REVOLUTION



PIONEERING IoT SINCE 2001

**TRI**  
Programmable Logic Controllers

Powerful & Easy Ladder  
+ BASIC Programming  
Ethernet integrated  
MODBUS TCP/IP  
DI/Os & AI/Os integrated

OEM Prices as low as \$119  
for full-feature Nano-10 PLC

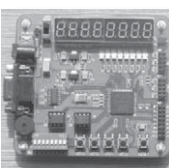
tel : 1 877 TRI-PLCS  
web : [www.triplc.com/cci.htm](http://www.triplc.com/cci.htm)



**TRI** TRIANGLE  
RESEARCH  
INTERNATIONAL

**LEARN FPGA PROGRAMMING  
USING THE ALTERA  
QUARTUS II AND VERILOG**

The course comes with a Cyclone board, as shown, and an USB Blaster. Also included is a 175-page manual containing ten chapters that are full of examples and tutorials. This in-depth course is a must for the hobbyist, as well as, for schools teaching FPGAs. The manual is easy to understand and correct.



\$55 + \$12 shipping

**ADVANCED MICRO CIRCUITS**  
443-817-2557  
[www.amcmicros.com](http://www.amcmicros.com)  
[information@amcmicros.com](mailto:information@amcmicros.com)

**MaxBotix**<sup>®</sup>  
High Performance Ultrasonic Rangefinders

**Save 10%**  
Online Web-Order Code:  
**C140731**  
Valid thru August 1, 2014

**HRXL-MaxSonar<sup>®</sup>-WR<sup>™</sup>**

- High noise tolerance
- IP67 rated
- 1 mm resolution
- Multi-Sensor operation
- Calibrated beam pattern
- Starting at \$109.95



**XL-MaxSonar<sup>®</sup>-EZ<sup>™</sup>**

- Great for UAV's and robotics
- Incredible noise immunity
- Small in size
- 1cm resolution
- Automatic calibration
- Starting at \$39.95



Phone: 218-454-0766 Email: [sales@maxbotix.com](mailto:sales@maxbotix.com)  
[www.maxbotix.com](http://www.maxbotix.com)

**Real Logic Analyzers** from **\$349.00**

with Decoders & Plug-in-Kit included



- 9 to 36 Channels
- 100 Mps to 400 Mps
- Store & Forward Architecture
- Real-time Smart Compression
- Plot & Measurement Features
- Powerfull Search Tools
- Plug-ins for full Custom Protocol integration

by **TechTools** [www.tech-tools.com/real.htm](http://www.tech-tools.com/real.htm) • (972) 272-9392

**Connect**  
with engineers  
seeking your  
design solutions.


Magazine, e-newsletters,  
website, advertorials, and more.

**RESERVE ADVERTISING  
SPACE TODAY!**

Strategic Media Marketing, LLC  
**978.281.7708**  
[circuitcellar@smmarketing.us](mailto:circuitcellar@smmarketing.us)  
[www.smmarketing.us](http://www.smmarketing.us)

**microEngineering Labs, Inc.**  
[www.melabs.com](http://www.melabs.com) 888-316-1753

**Programmers for Microchip PIC<sup>®</sup> Microcontrollers**



**PC-Tethered USB Model (shown):**

- Standalone software
- Command-line operation
- Hide GUI for automated use
- Override configuration with drop-downs

**Stand-Alone Field Programmer:**

- Power from target device or adapter
- Program file stored on SD-CARD
- Programming options stored in file
- Single-button operation

**Starting at \$79.95**

Program in-circuit or use adapters for unmounted chips.  
Zero-Insertion-Force Adapters available for DIP, SOIC, SSOP, TQFP, and more.  
PIC is a registered trademark of Microchip Technology Inc. in the USA and other countries.

# Wearable Medical Computing and the Amulet Project

By Vivian Genaro Motti



Dr. Vivian Genaro Motti (vgenaro@clemsun.edu) holds a PhD in Human Computer Interaction from the Université catholique de Louvain in Belgium, is a Postdoctoral Research Fellow in the School of Computing at Clemson University in Clemson, SC. She works on the Amulet project ([www.amulet-project.org](http://www.amulet-project.org)), which is funded by a three-year, \$1.5 million grant from the National Science Foundation's Computer Systems Research program. Vivian investigates wearability and privacy in wearable applications for mobile health. Vivian has a BA in Biomedical Informatics and an MS in Human Computer Interaction from University of Sao Paulo in Brazil.



The Amulet bracelet-style prototype for developers enables users to control its specific settings, for instance, by replacing its board.

Health care is one of the most promising areas for employing wearable devices. Wearable mobile health (mHealth) sensors can track activities (e.g., count steps or calories), monitor vital signs (e.g., heart rate and blood pressure), measure biometric data (e.g., glucose levels and weight), and provide alerts about medical emergencies (e.g., heart failures, falls, and shocks).

Applying wearable computing to support mHealth is promising but involves significant risks. They include the security issues related to the reliability of the sensors employed, the accuracy of the data collected, and the privacy of sensitive information.

Under the federally funded Amulet project, an interdisciplinary team of Dartmouth College and Clemson University researchers is investigating how wearable devices can effectively address medical problems while ensuring usability and security for mHealth applications.

The project aims to develop pieces of Amulet "computational jewelry," which are wearable mHealth devices, and a software framework for monitoring them. A discreet Amulet pendant or bracelet would interact with other wearable health sensors that constitute the wearer's wireless body-area network (WBAN). The Amulet device would serve as a "hub" that tracks health information from wearable sensors and securely sends it to other devices.

The project's goals are multifold. Regarding the hardware, we're focusing on designing small and unobtrusive form factors with sensing capabilities and efficient power sources. With respect to the software, we are concentrating on processing and interpreting the digital signs coming from the sensors, effectively communicating and synchronizing data with external devices, and managing encryption.

Amulet's hardware architecture includes an application processor to perform computationally intensive tasks and a coprocessor to manage radio communications and internal sensors. Amulet's current prototypes, which are wearable as wristbands, contain an accelerometer and a gyroscope to monitor the wearer's activities, a magnetometer, a temperature sensor, a light sensor, and a microphone. To save power, the application processor is powered off most of the time while the coprocessor handles all real-time device interactions.

Using an event-driven software architecture, Amulet enables applications to survive routine processor shutdowns. Amulet is reactive, running only when an interesting event occurs.


To handle such events, programmers define their application as a finite-state machine with appropriate event-handler functions. Amulet's architecture enables applications to store information to be retained between events or to append it to a data log. Explicitly managing a program state (rather than implicitly managing a state in a thread's run-time stack) enables the run-time system to save the application state to persistent memory and power off the main processor without harming applications.

Amulet provides a secure solution that ensures the accuracy and integrity of sensed and transmitted data, continuous service availability (e.g., data sensing and processing and sending alerts and notifications), and access to the device's data and services only by authorized parties after their successful authentication.

Two key features help Amulet provide security in mHealth applications: code sandboxing, which enforces access control, protects memory, and restricts the execution of the application's event handlers; and the authorization manager, which verifies all access policies and service requests before forwarding them to a corresponding service manager.

To protect privacy, Amulet's devices enable users to control what is sensed and stored, where it is stored, and how it is shared (with whom). By employing privacy policies, we intend to protect patients' sensitive information and ensure confidentiality through authorized access and controlled sharing.

Amulet aims to create intuitive interfaces and multimodal interaction for user input and output (e.g., using gestures and haptics). To guarantee easy wearability, the Amulet team focuses on understanding the user's wishes, needs, and requirements and translating them into appropriate design decisions. The team has developed a list of principles and guidelines for wearability, which will help designers ensure high levels of comfort, aesthetics, ergonomics, and discretion in their projects.

With a general-purpose solution for body-area mHealth, Amulet complements the capabilities of a smartphone and facilitates the development of applications that integrate one or more mHealth wearable devices. Amulet solutions envisage different applications, including support for emergency response (e.g., providing immediate notifications and quick responses in medical emergencies), stress monitoring, smoking cessation, diet (e.g., bite counting), and physical therapy for knee injury recovery. 

# INTRODUCING CC Vault

The newest archive product from your friends at *Circuit Cellar*!



This pocket-sized vault comes fully loaded with every issue of *Circuit Cellar* magazine and serves as an unparalleled resource for embedded hardware and software design tips, schematics, and source code.

From green energy design to 'Net-enabled devices, maximizing power to minimizing footprint, CC Vault\* contains all the trade secrets you need to become a better, more educated electronics engineer.

**BONUS!** Build your archive by downloading your latest-issue PDFs straight to the drive! Personalize your CC Vault by adding *Elektor* or *audioXpress* issue archives, available as an add-on during time of purchase, or your very own project files.

\*CC Vault is a 16-GB USB drive.



## Order yours today! [cc-webshop.com](http://cc-webshop.com)

PRINTED CIRCUIT BOARDS & ASSEMBLY

THERE  
HAS TO BE A  
CATCH!

**FREE  
PARTS?**

**No Catch. This is for Real**  
No Quantity Restriction

**LET US FABRICATE AND ASSEMBLE YOUR BOARDS,  
AND WE WILL PAY UP TO \$500 FOR THE PARTS.**

FIRST TIME CUSTOMERS ONLY



**imagineering inc.**

[www.PCBnet.com](http://www.PCBnet.com) • (847) 806-0003 • [sales@pcbnet.com](mailto:sales@pcbnet.com)