



circuit cellar

INTERNET OF THINGS INNOVATION


Build a Microcontroller-Based Monitoring System

Base Unit Design, Sensor
Unit Integration, Software
Development, & More

- Q&A: Cutting-Edge Haptics Research | Editors' Picks: Robotics Projects
- Plug-In Sweep Generator Board | How to Mimic IR Control Signals
- | The Sentry Monitoring Project | Sensorless Maximum Power Point Tracking
- Cooling Awareness for HPC | Achieve Electromagnetic Compliance
- | Handle Startup Transients | More Solid-State Lighting Tips
- | E-Compassing for Mobile Robotics ■ The Future of Flexible Circuits




Ethernet Core Modules with High-Performance Connectivity Options



- **MOD5270**
147.5 MHz processor with 512KB Flash & 8MB RAM · 47 GPIO · 3 UARTs · I²C · SPI
- **MOD5234**
147.5 MHz processor with 2MB flash & 8MB RAM · 49 GPIO · 3 UARTs · I²C · 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 I²C · 3 SPI · 2 CAN · SSI · 8 ADC · 2 DAC · 8 PWM · 1-Wire[®] interface
- **NANO54415**
250 MHz processor with 8MB flash & 64MB RAM · 30 GPIO · 8 UARTs · 4 I²C · 3 SPI · 2 CAN · SSI · 6 ADC · 2 DAC · 8 PWM · 1-Wire[®] 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.

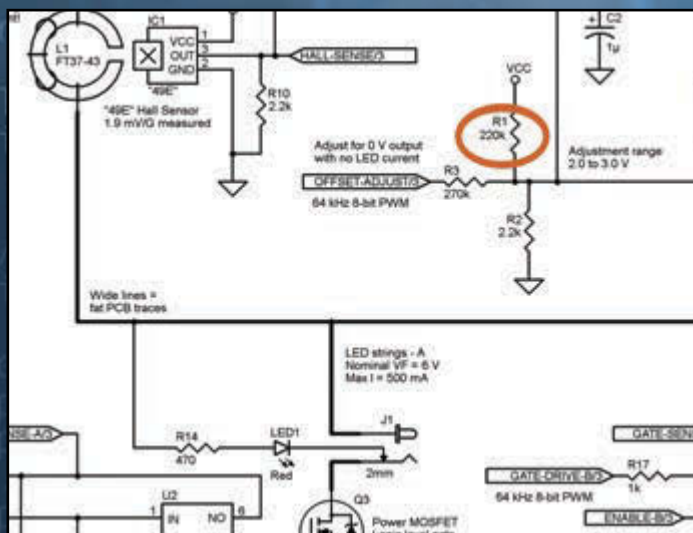
➤ **For additional information please visit**
<http://www.netburner.com/kits>



MONTHLY ENGINEERING CHALLENGE

Each month, you're challenged to find an error in a schematic or in code that's presented on the challenge webpage. Locate the error for a chance to win prizes and recognition in Circuit Cellar magazine!

Prizes such as a NetBurner MOD54415 LC Development kit or a Circuit Cellar subscription will be announced each month.



```

1 #include <stdio.h>
2
3 int main()
4 {
5     int first, second, sum;
6     int *p, *q;
7
8     printf("Enter two integers to add: \n");
9     scanf("%d%d", &first, &second);
10
11     p = &first;
12     q = &second;
13
14     sum = p + q;
15
16     printf("Sum of entered numbers = %d\n", sum);
17
18     return 0;
19 }

```

Participate: circuitcellar.com/engineering-challenge-netburner

Launch: 1st of each month

Deadline: 20th of each month

No purchase necessary to enter or win. Void where prohibited by law. Registration required. Prizes subject to change based on availability. Review these terms before submitting each Entry. More info: circuitcellar.com/engineering-challenge-netburner-terms

Issue 296 March 2015 | 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 © 2015 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 infringement liability for constructing or operating such devices.

© Circuit Cellar 2015 Printed in the United States

TIMELY TOPICS: ROBOTICS, WEARABLES, & MORE

In this issue we cover some of most important subjects of our time: robot systems and their use in innovative applications such as surgical tools; human- and building-monitoring systems; machine learning; energy efficiency; cutting-edge lighting; flexible circuitry; and more. Here are the details.

On page 6, we present an interview with University of Pennsylvania professor Katherine Kuchenbecker, who is helping lead a revolution in haptics technology. Her goal: "We should be able to capture and reproduce how surfaces feel with the same acuity that we can capture and reproduce how surfaces look."

Turn to page 12 to learn about four of our staff's favorite robotics projects: an inertial rolling robot, a photo-taking mobile robot, a robot navigation system, and compact "minirobot."

Gerard Fonte wraps up his two-part series, "Electronics Testing Platform," with details about a custom sweep generator board (p. 16). This project is a plug-in board for the base unit he covered last month.

Looking to add infrared control capability to an electronic system? Check out Tommy Tyler's article, "Infrared Learning," in which he presents a low-cost method for mimicking IR control signals (p. 26).

David Penrose's Sentry Project on page 32 won Honorable Mention in the WIZnet Connect the Magic 2014 Design Challenge. Follow his lead to construct a building-monitoring system of your own.

Maximum power point trackers (MPPTs) ensure that maximum power is transferred to a device from renewable energy input applications. In "Sensorless Maximum Power Point Tracking," you'll learn about a new maximum power tracking control method (p. 38).

High-performance computing (HPC) clusters run intensive, multi-processor applications. On page 46, Ayse Coskun discusses how cooling awareness in job allocation can improve energy efficiency in such environments.

In February, George Novacek discussed practical EMC requirements for electronic systems. This month, he finishes his multipart series with tips for achieving electromagnetic compliance.

On page 54 Ed Nisley details several parts of a recent machine controller project that were unexpectedly problematic. Use his "debugging guide" when you're faced with project issues of your own.

Last month, Jeff Bachiochi introduced the topic of electroluminescent (EL) technology. In this issue he explains how EL technology can produce a nondirectional output as opposed to an LED's concentrated point source output (p. 62).

Looking for an interesting robotics project? We pulled Joseph Miller's "Mini Rover 7" article out of the archives to get you motivated to begin a new project (p. 68).

Mark Finstad wraps up the issue with an essay on the future of flexible circuitry (p. 80). He prognosticates that they future of flex circuit technology is "on you"—that is, in wearable technology.

C. J. Abate

cabate@circuitcellar.com

THE TEAM

EDITOR-IN-CHIEF

C. J. Abate

ART DIRECTOR

KC Prescott

ADVERTISING COORDINATOR

Kim Hopkins

PRESIDENT

Hugo Van haecke

COLUMNISTS

Jeff Bachiochi (From the Bench), Ayse K. Coskun

(Green Computing), Bob

Japenga (Embedded

in Thin Slices), Robert

Lacoste (The Darker

Side), Ed Nisley (Above

the Ground Plane),

George Novacek (The

Consummate Engineer),

and Colin O'Flynn

(Programmable Logic in Practice)

FOUNDER

Steve Ciarcia

PROJECT EDITORS

Chris Coulston, Ken

Davidson, and David

Tweed

OFFICE ASSISTANT

Alice Roberts

OUR NETWORK



SUPPORTING COMPANIES


Accutrace	7	Measurement Computing Corp.	25
All Electronics Corp.	79	microEngineering Labs	79
Anaren	15	MyRO Electronic Control Devices, Inc.	79
AP Circuits	53	NetBurner, Inc.	C2, 1
Custom Computer Services	79	Newhaven Display International	41
EMAC, Inc.	19	Pico Technology, Ltd.	31
IAR Systems	43	Saelig Co., Inc.	19
Imagineering, Inc.	C4	Technologic Systems	59, 73
Ironwood Electronics	79	Teledyne LeCroy Corp.	11
Jeffery Kerr, LLC	29	Triangle Research International, Inc.	79
MaxBotix, 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

MARCH 2015 • ISSUE 296

ROBOTICS

CUTTING-EDGE HAPTICS RESEARCH



STAFF PICKS: INNOVATIVE ROBOTICS PROJECTS

CC COMMUNITY**06 : QUESTIONS & ANSWERS****Advances in Haptics Research**

An Interview with Katherine J. Kuchenbecker

Details on the design and control of haptic interfaces for applications such as robot-assisted surgery

INDUSTRY & ENTERPRISE**12 : EDITORS' PICKS**

Four of the CC staff's favorite robotics projects

14 : PRODUCT NEWS**15 : CLIENT PROFILE**

Freescale Semiconductor (Costa Mesa, CA)

FEATURES**16 : Electronics Testing Platform (Part 2)**

DDS Sweep Generator

By Gerard Fonte

How to construct a plug-in sweep generator board

26 : Infrared Learning

How to Mimic IR Control Signals

By Tommy Tyler

A small, inexpensive module for learning and reproducing infrared signals

32 : The Sentry Project*By David Penrose*

Build a compact microcontroller-based, building-monitoring system


38 : Sensorless Maximum Power Point Tracking*By David Weight, Ian Johnson, and Sajjad H. Lalji*

An impedance-matching method for a closed-loop digital control system




A LOOK BACK AT THE MINI ROVER 7 ROBOT

CONTENTS



MCU-BASED SENTRY MONITORING SYSTEM



MAXIMUM POWER POINT TRACKING

68 : FROM THE ARCHIVES**Mini Rover 7***By Joseph Miller*

An overview of the Mini Rover 7 robot, which is modeled after the NASA/JPL Rocky 7 Mars rover

COLUMNS**46 : GREEN COMPUTING****Cooling Awareness in HPC Job Allocation***By Ayse K. Coskun*

How cooling awareness in job allocation can improve energy efficiency

50 : THE CONSUMMATE ENGINEER**Essential Electromagnetic Compliance (Part 2)**

Tips for Achieving Compliance

By George Novacek

Useful advice for dealing with EMI

54 : ABOVE THE GROUND PLANE**Startup Transients***By Ed Nisley*

Dealing with startup problems arising in a retrofitted sewing machine controller

62 : FROM THE BENCH**Solid-State Lighting (Part 2)**


Get Your Lighting Under Control

By Jeff Bachiochi

How EL technology can produce a nondirectional output

TESTS & CHALLENGES**77 : TEST YOUR EQ****78 : CROSSWORD****TECH THE FUTURE****80 : The Future of Flexible Circuitry***By Mark Finstad*

Why flexible circuitry will grow in popularity as the interconnect method of choice



TIPS FOR ELECTROMAGNETIC COMPLIANCE

QUESTIONS & ANSWERS

Advances in Haptics Research

An Interview with Katherine J. Kuchenbecker

Katherine J. Kuchenbecker is an Associate Professor in Mechanical Engineering and Applied Mechanics at the University of Pennsylvania, with a secondary appointment in Computer and Information Science. She directs the Penn Haptics Group, which is part of the General Robotics, Automation, Sensing, and Perception (GRASP) Laboratory. In this interview, she tells us about her research, which centers on the design and control of haptic interfaces for applications such as robot-assisted surgery, medical simulation, stroke rehabilitation, and personal computing.



CIRCUIT CELLAR: When did you first become interested in haptics and why did you decide to pursue it?

KATHERINE: I chose to become an engineer because I wanted to create technology that helps people. Several topics piqued my interest when I was pursuing my undergraduate degree in mechanical engineering at Stanford, including mechatronics, robotics, automotive engineering, product design, human-computer interaction, and medical devices. I was particularly excited about areas that involve human interaction with technology. Haptics is the perfect combination of these interests because it centers on human interaction with real, remote, or virtual objects, as well as robotic interaction with physical objects.

My first exposure to this field was a “haptic paddle” lab in a Stanford course on system dynamics, but that alone wouldn’t have been enough to make me fall in love with this field. Instead, it was conversations with Günter Niemeyer, the professor who advised me in my PhD at Stanford. I knew I wanted a doctorate so that I could become a faculty member myself, and I was inspired by the work he had done as an engineer at Intuitive Surgical, Inc., the maker of the da Vinci system for robotic surgery. Through my early research with Günter, I realized that it is incredibly satisfying to create computer-controlled electro-mechanical systems that enable the user to touch virtual objects or control a robot at a distance. I love demonstrating haptic systems because people make such great faces when they feel how the system responds to their movements.

Another great benefit of studying haptics is that I get to work on a wide variety of applications that could potentially impact people in the near future: robotic surgery, medical training, stroke rehabilitation, personal robotics, and personal computing, to name a few.

CIRCUIT CELLAR: What is haptography? What are its benefits?

KATHERINE: I coined the term “haptography” (haptic

photography) to proclaim an ambitious goal for haptics research: we should be able to capture and reproduce how surfaces feel with the same acuity that we can capture and reproduce how surfaces look.

When I entered the field of haptics in 2002, a lot of great research had been done on methods for letting a user feel a virtual three-dimensional shape through a stylus or thimble. Essentially, the user holds on to a handle attached to the end of a lightweight, back-drivable robot arm; the 3D Systems Touch device is the most recent haptic interface of this type. A computer measures the motion that the person makes and constantly outputs a three-dimensional force vector to give the user the illusion that they are touching the object shown on the screen. I was impressed with the haptic demonstrations I tried back in 2002, but I was also deeply disappointed with how the virtual surfaces felt. Everything was soft, squishy, and indistinct compared to how real objects feel. That’s one of the benefits of being new to a field; you’re not afraid to question the state of the art.

I started working to improve this situation as a doctoral student, helping invent a way to make hard virtual surfaces like wood and metal feel really hard and realistic. The key was understanding that the human haptic perceptual system keys in on transients instead of steady-state forces when judging hardness. I had to write a research statement to apply for faculty positions at the end of 2005, so I wrote all about haptography. Rather than trying to hand-program how various surfaces should feel, I wanted to make it all data driven. The idea is to use motion and force sensors to record everything a person feels when using a tool to touch a real surface. We then analyze the recorded data to make a model of how the surface responds when the tool moves in various ways. As with hardness, high-frequency vibration transients are also really important to human perception of texture, which is a big part of what makes different surfaces feel distinct. Standard haptic interfaces weren’t designed to output high-frequency vibrations, so we typically attach a voice-coil actuator (much

PRINTED CIRCUIT BOARDS

THINK YOU CAN FIND PCB PRICES THAT BEAT OURS?

WE DARE YOU.

If you do, than we
will match the price
AND give you \$100
towards your
next order!

THERE ARE NO GAMES INVOLVED IN OUR PRICING



Our Capabilities:

- 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), Polymide, 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 Microvia, Conductive Plate Shut.

Take the Accutrace Challenge and see WHY OUR PRICING CANNOT BE BEATEN

Accutraceinc.

www.PCB4u.com sales@PCB4u.com

QUESTIONS & ANSWERS

like an audio speaker) to the handle, near the user's fingertips. When the user is touching a virtual surface, we output data-driven tapping transients, friction forces, and texture vibrations to try to fool them into thinking they are touching the real surface from which the model was constructed. After many years of research by my PhD students Heather Culbertson and Joe Romano, we've been able to create the most realistic haptic surfaces in the world.

My work in haptography is motivated by a belief that there are myriad applications for highly realistic haptic virtual surfaces. One exciting use is in recording what doctors and other clinical practitioners feel as they use various tools to care for their patients, such as inserting an epidural needle or examining teeth for decay (more on this below). Haptography would enable us to accurately simulate those interactions so that trainees can practice critical perceptual-motor skills on a computer model instead of on a human patient. Another application that excites us is adding tactile feedback to online shopping. We'd love to use our technology to let consumers feel the fabrics and surfaces of products they're considering without having to visit a physical store. Touch-mediated interaction plays an important role in many facets of human life; I hope that my team's work on haptography will help bring highly realistic touch feedback into the digital domain.

CIRCUIT CELLAR: You founded the Penn Haptics Group in 2007. When you started the group, what was your mission?

KATHERINE: I wanted to create a world-class research group dedicated to inventing clever and useful haptic technologies. I started at the University of Pennsylvania (Penn) as an Assistant Professor, which means I had six years to prove that I was worthy of tenure (a job for life). While going through the tenure process can be stressful, I found it exhilarating to finally have the opportunity to assemble my own team of students and work on problems and ideas that interested us. Scientifically, I have focused on elevating and formalizing our understanding of haptic feedback while simultaneously uncovering new opportunities for its use in interactions between humans, computers, and machines. By working really hard, and through the hard work of all my great students, I was able to establish a research group that is now reasonably well respected in my field. We have been able to get research grants, publish papers, and even win some awards for our work. I'm proud of the work we've been able to do, and I'm equally proud of the students who have been working by my side. They taught me a lot!



This wearable device allows the user to control the motion of the PR2 robot's hand and also feel what the PR2 is feeling. The haptic feedback is delivered via a geared DC motor and two voice-coil actuators.

CIRCUIT CELLAR: How has the Haptics Group grown or evolved over the past few years? How many students, researchers, and professors are currently involved with the group?

KATHERINE: My group has both grown and diversified over the last eight years. I started out with just one PhD student, plus a bunch of students who took my graduate-level class on haptic interfaces. Back in those days, I was delighted and a little surprised to find that students wanted to learn more about what I study. I like to incorporate each student's interests and ideas into the research we do together, which tends to lead us in lots of interesting directions. For example, one of my current PhD students (Naomi Fitter) is really interested in using robots for education, so we've been working to create a robot that can teach people how to play hand-clapping games. My lab currently includes one postdoctoral researcher, five doctoral students, about eight masters students, and four undergraduates. These students hail from mechanical engineering, electrical engineering, computer

science, and bioengineering, and I've also worked with several medical students and surgical residents. Overall, I've advised about 40 graduate students and 50 undergraduates in research. I've also hosted visiting graduate students from Italy, Germany, Japan, and South Korea, and I had a prominent Korean haptics professor spend one year as a visiting faculty member in my lab. These days, I get a lot of emails from students who want to join my lab. I'm honored by this attention, but I am sad to have to turn away so many students want to work on the same kinds of topics that excite me.

CIRCUIT CELLAR: Which of the group's projects most interest you at this time?

KATHERINE: That's a hard question! I'm excited about all of the projects we are pursuing. There are a few I can't talk about, because we're planning to patent the underlying technology once we confirm that it works as well as we think it does. Two of those that are in the public domain have been fascinating me recently.

Tactile Teleoperation: My lab shares a Willow Garage PR2 (Personal Robot 2) humanoid robot with several of the other faculty in Penn's GRASP Lab. Our PR2's name is Graspy. While we've done lots of fun research to enable this robot to autonomously pick up and set down unknown objects, I'd always dreamed of having a great system for controlling Grasp from a distance. Instead of making the operator use a joystick or a keyboard, we wanted to let him or her control Grasp using natural hand motions and also feel what Grasp

QUESTIONS & ANSWERS

was feeling during interactions with objects. My PhD student Rebecca Pierce recently led the development of a wearable device that accomplishes exactly this goal. It uses a direct-drive geared DC motor with an optical encoder to actuate and sense a revolute joint that is aligned with the base joint of the operator's index finger. Opening and closing your hand opens and closes the robot's parallel-jaw gripper, and the motor resists the motion of your hand if the robot grabs onto something. We supplement this kinesthetic haptic feedback with tactile feedback delivered to the pads of the user's index finger and thumb. A voice coil actuator mounted in each location moves a platform into and out of contact with the finger to match what the robot's tactile sensors detect. Each voice coil presses with a force proportional to what the corresponding robot finger is feeling, and the voice coils also transmit the high-frequency vibrations (typically caused by collisions) that are sensed by the MEMS-based accelerometer embedded in the robot's hand. We track the movement of this wearable device using a Vicon optical motion tracking system, and Grasp follows the movements of the operator in real time. The operator sees a video of the interaction taking place. We're in the process of having human participants test this teleoperation setup right now, and I'm really excited to learn how the haptic feedback affects the operator's ability to control the robot.

Palpation for Robotic Surgery: My PhD student Will McMahan and I previously invented a new way to add haptic feedback to robotic surgery systems like the Intuitive da Vinci; we use tiny accelerometers to measure the high-frequency vibrations of the robotic instruments, and we use voice-coil actuators to play those vibrations for the surgeon to feel in real time. While we're still trying to convince a company to incorporate our invention into their robotic surgery system, we've continued to work on new ways of adding haptic feedback to robotic surgery. One big drawback to minimally invasive surgery is that the surgeon can't tactilely examine soft tissue for anomalies such as tumors. We have been addressing this problem in two distinct research thrusts, both of which use a SynTouch BioTac biomimetic tactile sensor to measure rich haptic data during interactions with simulated samples of soft tissue. In the first project, Claudio Pacchierotti (an Italian PhD student visiting my lab) worked with other students to attach the BioTac to the end of a da Vinci instrument, so we could use it to palpate silicone rubber tissue samples that contained simulated tumors. We then created a custom tactile feedback device that used three small RC servo motors and an ALPS ForceReactor vibration actuator to move a platform below the user's fingertip. We mounted this tactile device on the hand controller that the surgeon uses to control the instrument to which the BioTac was attached. We then wrote an algorithm that continually processes the finger deformations and vibrations that BioTac feels and drives the tactile device to output an approximation of these haptic signals for the operator to feel. It works really well! In the



A high-bandwidth MEMS-based accelerometer records the sensations a dentist feels as she probes an extracted human tooth. Feeling these recordings lets dental trainees practice diagnosing dental decay before they treat live patients.

second half of the palpation project, my PhD student Jennifer Hui has been writing algorithms to automatically determine whether a given sample of tissue contains a tumor. Aided by an undergraduate researcher, she mounted the BioTac on a CNC mill and recorded data while it methodically pressed into a large sample of simulated tissue. Now she is using machine learning principles to write software that can process what the BioTac feels in a given location and decide whether it feels like a tumor, the edge of a tumor, or healthy tissue. In the long run, we'd like to combine these two approaches so that a surgeon could control the tactile sensor, feel what it is feeling, and receive an automatic second opinion on whether that tissue contains a tumor.

CIRCUIT CELLAR: In your TEDYouth talk, you describe a project in which a dental tool is fitted with an accelerometer to record what a dentist feels and then replay it back for a dental student. Can you tell us a bit about the project? How it started and how it works?

KATHERINE: This project spun out of my haptography research, which I described above. While we were learning to record and model haptic data from interactions between tools and objects, we realized that the original recordings had value on their own, even before we distilled them into a virtual model of what the person was touching. One day I gave a lab tour to two faculty members from the Penn School of Dental Medicine who were interested in new technologies. I hit it off with Dr. Margrit Maggio, who had great experience in teaching general dentistry skills to dental students. She explained that some dental students really struggled to master some of the tactile judgments needed to practice dentistry, particularly in discerning whether or not a tooth surface is decayed (in popular parlance, whether it has a cavity). A few students and I went over to her lab to test whether our accelerometer-based technology could capture the subtle details of how decayed vs. healthy tooth tissue feels. While the recordings are a little creepy to feel, they are super accurate. We refined our approach and conducted several studies on the potential of this technology to be used in training dental students. The results were really encouraging, once again showing the potential that haptic technology holds for improving clinical training.

CIRCUIT CELLAR: You were a TA in the Stanford Product Realization Lab (PRL). What sort of work did you do there?

KATHERINE: When I was an undergrad, I thought the PRL was the coolest place on Stanford campus; it's a giant machine shop, rapid-prototyping lab, CAD loft, model shop, welding room, and foundry. The PRL is staffed by ubercool Shop TAs—engineering grad students tasked with teaching other students how to use all that fancy equipment without injuring themselves, each other, or the machines. Given my

QUESTIONS & ANSWERS

esteem for their profession, I applied to become a Shop TA as soon as I got into grad school, and I feel incredibly lucky that they hired me. Being a Shop TA was an amazing learning experience. I worked five four-hour shifts each week, sometimes teaching structured labs and other times running open shop sessions where I would be hounded by dozens of students who needed help getting their projects done. I got to teach manual milling, manual turning, CNC milling, oxy-acetylene welding, and sand casting; for this last one, I learned to run the furnace to melt and then pour molten bronze! My favorite time to be in the shop was the Friday night session, which ran from 7 pm to 11 pm. Working on so many different projects gave me great practice at mechanical design and problem solving, and it showed me that I absolutely love teaching and working with students. Being around such a creative and dedicated group of people was inspiring, as was having such good resources for manufacturing. During my time as a Shop TA, I designed and an organically shaped pendant lamp and CNC milled it out of aluminum and clear acrylic. I also worked with two other students and a urologic surgeon to create the NeoDriver, a surgical needle driver that allows the surgeon to re-angle the suture needle in the jaws of the tool with just the touch of a finger, without needing his or her other hand.

CIRCUIT CELLAR: Tell us about your postdoctoral research at Johns Hopkins.

KATHERINE: After I finished my PhD, I spent one year doing research with Professor Allison Okamura at the Johns Hopkins University. She is a worldwide expert on haptics, and I loved seeing the fascinating range of projects she and her lab were tackling—everything from human motor control to steerable needles.

The main focus of my work that year was related to the Revolutionizing Prosthetics program funded by DARPA. Many people in haptics agreed that next-generation upper-limb prostheses should provide the wearer with haptic feedback such as grip force or the making and breaking of contact. Thinking farther ahead, we were curious about whether prosthetic devices should give the wearer an artificial sense of proprioception, which is how you can feel the position and orientation of your limbs when you close your eyes.

Proprioception is known to be important in human motor control, but no one had studied whether it might help humans control artificial limbs. Because we wanted to compare perfect proprioception with no proprioception, I created an apparatus that simulated this scenario by giving a user with an intact limb control over a virtual prosthetic device. We found that proprioceptive feedback did indeed seem to help users! In addition to this main project, I supervised three students in creating a tactile fingertip device that was an early pre-cursor to the tactile teleoperation projects that I discussed above.



This NeoDriver prototype lets the user change the angle of the suture needle in the jaws of the tool simply by adjusting the angle of a knob; stainless steel wires run from the knob along the left and right sides to rotate a tiny disk embedded in the jaws.


While it went very quickly, this year of postdoctoral research was a lovely oasis between the stress of finishing grad school and the strain of starting my lab at Penn. My favorite takeaways from Johns Hopkins were the friendships I made there, and I am especially grateful for Allison's mentorship. She's now a professor at Stanford, and two of my former students are working in her lab—what fun!

CIRCUIT CELLAR: Your work is becoming increasingly publicized. You have the TEDYouth talk, and the Haptics Group continues to grow. Do you frequently get requests to design solutions to problems? Are you frequently pitched product ideas or business opportunities? For instance, are medical teams, virtual reality startups, and other groups contacting you for prototyping and consulting work?

KATHERINE: I couldn't believe the amount of attention I got after my TEDYouth talk about haptics went live on TED.com. I initially tried to send a personal response to every email, but that quickly got to be too much. I am grateful that people do still contact me about interesting projects after finding my TED talk. For example, my students and I have recently been doing a project with researchers from a large European company; if I told you about this project, I'd have to kill you (as my Dad would say), but I can say that working on a totally new area has been intellectually invigorating. Most inquiries from startups don't go very far, but I recently agreed to work with Steven Domenikos, an experienced entrepreneur from Boston, to try to create a prototype for a new commercial device related to haptics. Again, I can't give any details, but it is fun to try to create something that would really touch the masses.

I do encourage people to contact me with any ideas and questions they have about haptics. My email address is kuchenbe@seas.upenn.edu I'll try to send you a personal response, but please don't be offended if that takes a few weeks!

CIRCUIT CELLAR: What is the "next big thing" in the field of haptics? Is there a specific area or technology that you think will be a game changer?

KATHERINE: Of course this depends on where you're looking. While cell phones and game controllers have had vibration alerts for a long time, I think we're just starting to see high-quality haptic feedback emerge in consumer products. Haptics can definitely improve the user experience, which will give haptic products a market advantage, but their cost and implementation complexity need to be low enough to keep the product competitive. On the research side, I'm seeing a big move toward tactile feedback and wearable devices. Luckily there are enough interesting open research questions to keep my students and me busy for 30 more years, if not longer! 

A new **wave** of thinking



WaveSurfer 3000
with **MAUI**



 **TELEDYNE LECROY**
Everywhere you look™

teledynelecroy.com/newwave


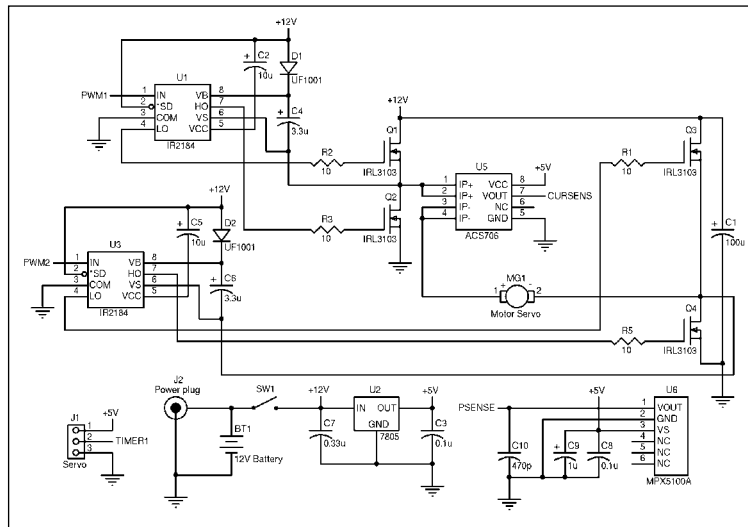
EDITORS' PICKS

Robotics

INERTIAL ROLLING ROBOT

Jeff Bingham & Lee Magnusson
Circuit Cellar 200, 2007

This H8/3664-based rolling robot is capable of inertial movement. A DC electric motor is attached to a pendulum and suspended inside an inflated ball, which provides the driving force. As the engineers explain, "The robot employs a single driven pendulum inside of a sphere. With 2 degrees of freedom, it drives and tilts the ball. With this geometry, the robot can reach high speeds, steer, and jump. In addition, it doesn't get stuck when it flips over."



The robot's inner core consists of an electronics assembly and a mechanical drive assembly. The spherical shell is an inflatable ball. The inner core is suspended via a drive axle from the two plastic caps.

The H8/3664 microcontroller connects to several components that provide control and sensing. The circuit is basic and easy to connect to most other microcontrollers as well.

These articles and others on topics relating to Robotics are available in the CC Webshop. Go to www.cc-webshop.com.

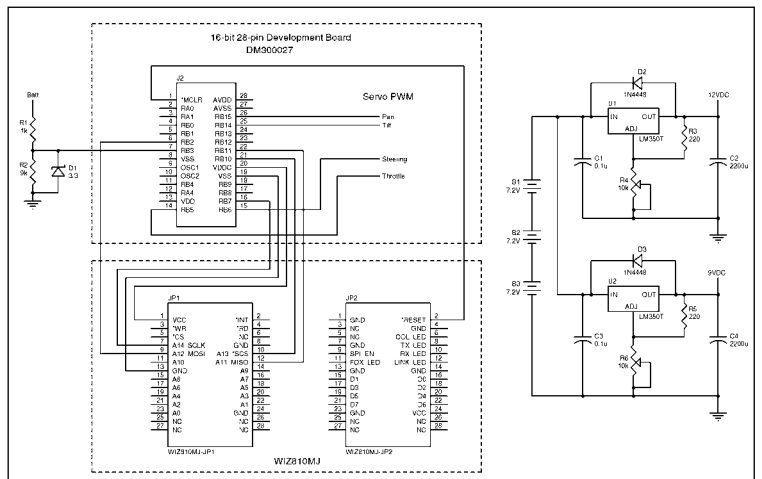
WIRELESS MOBILE ROBOTICS

Scott Coppersmith
Circuit Cellar 224, 2009

Scott used a microcontroller, an embedded Ethernet board, and a wireless router in an innovative control system for a compact mobile robot. The robot features a mounted webcam that transmits real-time pictures to a remote laptop. Scott explains how he planned the project, assembled the pieces, and created the control software.



The assembled robot with a webcam



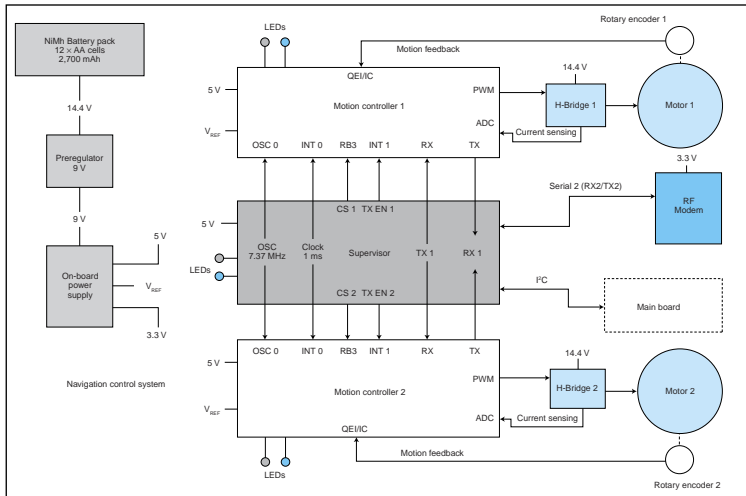
The embedded Ethernet board, speed control, and servos for controlling steering, camera pan, and tilt are wired to a Microchip Technology DM300027 16-bit, 28-pin development board. A WIZnet board is connected to the DM300027 development board by wire wrapping the header connector pins directly. The two boards communicate via the SPI bus.

EDITORS' PICKS


ROBOT NAV & CONTROL SUBSYSTEM

Guido Ottaviani
Circuit Cellar 224, 2009


You can build navigation control subsystem for an autonomous differential steering explorer robot. The system's software manages wheel speed, closed-loop control with a PID algorithm, dead-reckoning by odometry (in both theoretical and practical forms), field mapping, navigation, motor control, and more. The subsystem includes Guido's dsPIC-based navigation control board (the "dsNavCon") and an STMicroelectronics L298-based dual H-bridge board for controlling the geared 12-V motors.



The navigation control subsystem's main components



The entire navigation control subsystem



The motors, encoders, and H-bridges. The metal box at the top contains a battery pack, 9-V preregulators, and PI filters. It also acts as a heatsink.

DESIGN & PROGRAM A "MINIROBOT"

Marco Aiello
Circuit Cellar 236, 2010

Control capability is at the heart of most robotics applications, particularly when object avoidance is a project requirement. This mobile "minirobot" features a custom power board built around a Microchip Technology dsPIC30F2010 that translates serial commands into the right PWM power signals to drive the two motors and an OLIMEX SBC running a Debian distribution on an Atmel AT91SAM9260 ARM9 CPU.



A close-up shot of the motor board



The fully assembled Linux-based robot

COMMUNITY

PRODUCT NEWS

CONTROLLER FOR ONE-STEP DESIGN-IN OF 3-D GESTURE RECOGNITION

Microchip Technology recently announced a new addition to its patented GestIC family. The new MGC3030 3-D gesture controller features simplified user-interface options focused on gesture detection, enabling true one-step design-in of 3-D gesture recognition in consumer and embedded devices. Housed in an easy-to-manufacture SSOP28 package, the MGC3030 expands the use of 3-D gesture control features to high-volume, cost-sensitive applications such as audio, lighting, and toys.

The simplicity of gesture-detection integration offered by the MGC3030 is also achieved through Microchip's free, downloadable AUREA graphical user interface (GUI) and easily configurable general-purpose IO ports that even allow for host MCU/processor-free usage. The MGC3030's on-chip 32-bit digital signal processor executes real-time gesture processing, which eliminates the need for external cameras or controllers for host processing and allows for faster and more natural user interaction with devices..

The MGC3030 featuring GestIC technology is available in a 28-pin SSOP package. Each unit costs under \$2 each in high volumes.

Microchip Technology
www.microchip.com

CLAMSHELL SPRING PIN QFN SOCKET

Ironwood Electronics recently introduced a new QFN socket addressing high performance requirements for testing QFN devices - CBT-QFN-7039. The contactor is a stamped spring pin with 17-g actuation force per ball and cycle life of 50,000 insertions. The self inductance of the contactor is 0.75 nH, insertion loss less than 1 dB at 31.7 GHz. The current capacity of each contactor is 1.5 A at 20°C temperature rise. Socket temperature range is -55°C to 180°C.

The socket also features an alignment guide for precise device to pin alignment. The specific configuration of the package to be tested in the CBT-QFN-7039 is a QFN, 3 × 3 mm, 0.5-mm pitch, 10 positions with center ground pad. The socket is mounted using supplied hardware on the target PCB with no soldering. To use, place the QFN device into the socket base and lock the double latch socket lid on to the base using the latch.

The socket uses a compression wave spring to apply constant downward pressure enabling the device be interconnected to the target PCB. This socket can be used for hand test and quick device screening applications with the most stringent requirements.

The CBT-QFN-7039 costs \$552 with reduced pricing available depending on quantity required.

The CBT-QFN-7039 costs \$552 with reduced pricing available depending on quantity required.

Ironwood Electronics
www.ironwoodelectronics.com

SMALL STEP-DOWN POWER MODULE

Intersil Corp.'s ISL8203M is a dual 3A/single 6A step-down DC/DC power module that simplifies power supply design for FPGAs, ASICs, microprocessors, DSPs, and other point of load conversions in communications, test and measurement, and industrial systems. The module's compact 9.0 mm × 6.5 mm × 1.83 mm footprint combined with industry-leading 95% efficiency provides power system designers with a high-performance, easy-to-use solution for low-power, low-voltage applications.

The ISL8203M is a complete power system in an encapsulated module that includes a PWM controller, synchronous switching MOSFETs, inductors and passive components to build a power supply supporting an input voltage range of 2.85 to 6 V. With an adjustable output voltage between 0.8 and 5 V, you can use one device to build a single 6-A or dual output 3-A power supply.

Features and specifications: dual 3-A or single 6-A switching power supply; high efficiency, up to 95%; wide input voltage range: 2.85 to 6 V; adjustable output range: 0.8 to 5 V; internal digital soft-start: 1.5 ms; external synchronization up to 4 MHz; and overcurrent protection.

The ISL8203M power module costs \$5.97 in 1,000-piece quantities. The ISL8203MEVAL2Z evaluation costs \$67.

Intersil Corp.
www.intersil.com



FREQUENCY-PROGRAMMABLE TRANSCEIVER

The new Multi-Use Radio Service (Mini-MURS) NiM1B-154.570-5-12.5-MURS is a frequency-programmable narrow band transceiver that offers a low-power, reliable data link in a Lemos /Radiometrix transceiver standard pinout and footprint. It's suitable for licensed and unlicensed VHF allocations, FCC part 90 and part 95.

The transceiver's features include the following: conforms to EN 300 220-3 and EN 301 489-3 (10-mW version only); compliant with FCC part 95 (MURS); standard frequency 154.570 MHz or 154.600 MHz (reprogrammable); data rates up to 5 kbps for standard module; usable range over 1 km; feature-rich interface (true analog and/or digital baseband); and more.

Technical specifications:

- Fully integrated sigma-delta PLL synthesizer based design
- High stability TCXO reference
- Data bit rate: 5 kbps max.
- Transmit power: +13 dBm (20 mW)
- Image rejection: greater than 70 dB
- Receiver sensitivity: -120 dBm (for 12-dB SINAD)
- RSSI output with greater than 50-dBm range
- Supply: 3.3 V – 15 V at 30 mA transmit, 18 mA receive
- Dimensions: 33 × 23 × 11 mm (fully screened)
- Evaluation platforms: NBEK + BiM/SMX carrier

Lemos International
www.lemosint.com



Freescale Semiconductor

2100 East Elliot Road, Tempe, AZ 85284

www.freescale.com/Sensors

CONTACT: Freescale.com/Support

FEATURED PRODUCTS: Freescale is a global semiconductor company enabling the Internet of Tomorrow. As a leader in processing and sensing solutions, we are driving a more innovative and connected world for the future. Freescale sensors, smart digital networks and microcontrollers provide complete solutions for a secure Internet of Tomorrow. Expanding on its more than 35-year heritage of sensor innovation, Freescale is the only company supplying Internet of Tomorrow solutions from sensors (edge) to network to cloud.

WHY SHOULD CC READERS BE INTERESTED? Freescale's new FXAS21002C gyroscope can combine with accelerometers and magnetometers for a full 9-axis sensor fusion solution. The gyroscope has best-in-class power performance at 2.7 mA (active) and 2 μ A (standby) for advanced motion sensing applications. The FXAS21002C gyroscope has a dynamically selectable full scale range that is configurable for maximum angular velocity that the gyro can read allowing for multiple use cases to be addressed with a single device. For more information, please visit freescale.com/gyroscope.

CLIENT PROFILE



SPECIAL OFFER: Take your sensor design to the next level! Receive 50% off on the Freescale Development Platform for 9-Axis Sensor Solution. Visit freescale.com and search for kit number FRDM-STBC-AGM01. Select "Buy Direct" and enter code FSLND425 once the kit is in your shopping basket. One kit per person. Discount coupon valid only for purchases by persons 18 years of age or older made online at Freescale Buy Direct at freescale.com on or before December 31, 2015.



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

JOIN THE EVOLUTION.



Learn more



1905



1945



2005



Today

Get "mobile smart" in 3 easy steps:

- 1 Get your AIR for Wiced Smart dev kit at your distributor of choice. (See our website for a current list.)
- 2 Develop your wireless link and basic app using our exclusive Atmosphere development tool.
- 3 With our AIR for Wiced Smart module on board, proceed in record time to a prototype and final, mobile-app development!



Evolve to app-based control with AIR for Wiced Smart!

If you're ready to evolve from fixed control panels populated with dials, buttons, keypads, and LCD displays to mobile-app based control of your embedded product - check out Anaren's AIR for Wiced Smart module, featuring Broadcom's Wiced Smart Bluetooth® chip (BCM20737). Not only does our small-footprint, SMT, and pre-certified all-in-one module save you the time, effort, and trouble of designing your own radio... It's supported by our industry-exclusive Atmosphere development ecosystem that lets you develop your basic embedded code and app code in one, easy-to-use development tool - for a far speedier product development cycle and time-to-market.

Follow the steps at left to join the evolution, right now!

www.anaren.com/AIRforWiced
800-411-6596
In Europe: 44-2392-232392

Anaren®
What'll we think of next?™




Electronics Testing Platform (Part 2)

FEATURES

DDS Sweep Generator

In the first part of this series, Gerard explains how he engineered the base unit platform for the development of test instruments. Here he details the construction of a plug-in sweep generator board.

By Gerard Fonte (US)



If you've never used a sweep generator before, you don't know what you are missing. It's a very powerful test instrument that greatly simplifies and speeds up all sorts of analyses. You can quickly test the frequency response of an amplifier, speaker, or filter, find the precise resonant frequency of a circuit, determine just how long it takes a phase-lock-loop circuit to lock, and lots more. There are a number of inexpensive analog sweep generators available and they work acceptable well. However, this is a direct digital synthesis (DDS) sweep generator with features that can't be matched by any analog sweeper. Comparable DDS systems cost \$500 to \$600, many times more than this one.

This project is a plug-in board for the

base unit I described in the first part of this article series. I built the board around an inexpensive Analog Devices AD5930 DDS sweep generator chip (\$14 each). The AD5930 is a very powerful and flexible device with very nice specifications. It operates from 3 Hz to 25 MHz (Nyquist limit), although a more reasonable upper limit is 10 to 15 MHz. It has differential outputs (two outputs 180° out of phase with each other). Output distortion is typically about 0.1%. It can operate in fixed-frequency, burst, or continuous mode, linear or log sweep, with sine/triangle wave output as well as a separate square wave output (TTL level). Frequency step size is about 0.3 Hz for low frequency operation and 3 Hz for full-range operation. Frequency accuracy is

50 PPM. A software marker/cursor is provided that identifies a precise frequency sweep-point on your oscilloscope is extremely useful. Up to eight setups can be stored on EEPROM.

AMPLIFIER DESIGN APPROACH

The AD5930 does not provide much output drive capability. The peak-to-peak (PP) output voltage is only about 0.5 V and it can only drive with about 3 mA. Obviously this is not sufficient for most applications and an amplifier must be added.

There were conflicting goals in this amplifier design. The first was to provide a

fully integrated system for audio and other relatively low-frequency work (less than 1 MHz). This meant driving an 8-Ω speaker directly, small frequency steps (0.3 Hz), and software level controls.

The second goal was to provide the maximum frequency range available for RF applications (25 MHz/Nyquist). I wanted a high-level output (± 10 V unloaded) and capable of driving the standard 50-Ω loads. It also had to be stable driving any type of load.

Since I don't like choosing which is best, I provided both amplifiers (see **Figure 1**). The low-frequency amplifier uses a simple

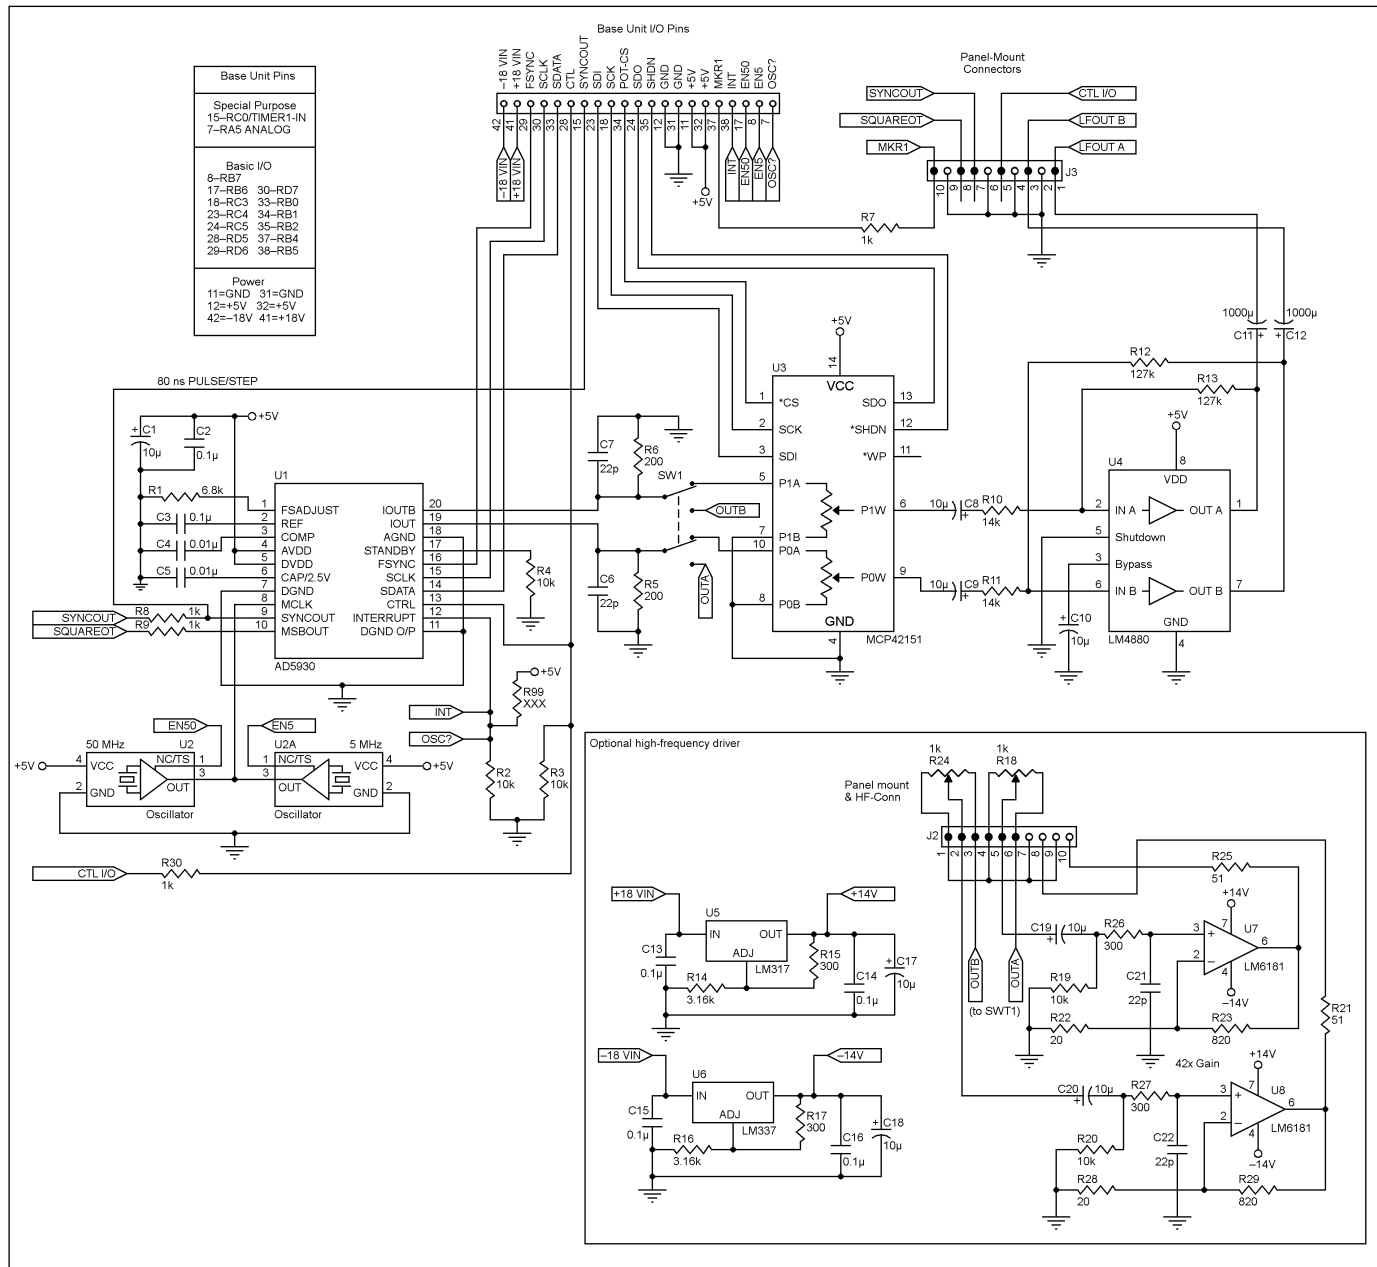



FIGURE 1 This is the sweep generator board with low-frequency and high-frequency amplifiers.

PHOTO 1


Jumpers can be used instead of a DPDT switch for single amplifier operation.



5-V, 0.25-W stereo audio amplifier (LM4880) to drive both channels. It operates up to 1 MHz, but not very well. I used a Microchip Technology MCP42151 digital potentiometer as a volume control. Initially, I wanted to use this for amplitude modulation (AM). However, this turned out to be impractical. The frequency

steps are about 0.3 Hz (using a 5-MHz clock).

The high-frequency amplifier uses a National Semiconductor LM6181 current feedback op-amp. (The typical op-amps found in most hobbyist designs are voltage feedback op-amps.) Current feedback (CF) op-amps have the very nice feature that the

**PHOTO 2**

J1 is mounted on the bottom of the PCB and consists of five pieces of SIP socket strip.

gain is nearly independent of the bandwidth. In this way, a power amplifier with a gain of 42 and maximum operating frequency of 25 MHz can be designed around a single amplifier stage. (A voltage feedback op-amp would need to have a gain-bandwidth product of over 1 GHz. While a few such devices were found none could drive 50 ohms.) A positive and negative 14-V supply is also required for this amplifier. Heatsinks are shown and listed in the parts list for the regulators but are not really required. They run cool.

There are a couple of important points to note about this CF circuit. The feedback resistor (R23 or R29) is specified and fixed at 820 Ω by the manufacturer. To change the gain, resistor R22 or R28 can be changed as required. The resistor-capacitor (RC) network (R26 + C21 or R27+C22) is used to stabilize the op-amp by reducing the high-frequency gain. The 51- Ω output resistor (R21 or R25) does double duty. The first is to isolate the op-amp from capacitive loads which could lead to instability. The second is to limit the output current in case of a short circuit.

CF op-amps tend to run hot, so a heatsink (with silicone grease) is an absolute requirement. Running into an open circuit (no load), the amp outputs 22 V_{pp} and warms up to




PHOTO 3

The TSSOP footprint for U1 is expanded for an adapter.

BEST SCOPES, LOWEST PRICES



PASSPORT-SIZE PC SCOPES \$162+

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/function gen. **PS2200A series**



30MHz SCOPE \$289

Remarkable 30MHz, 2-ch, 250MS/s sample rate scope. 8-in color TFT-LCD and AutoScale function. Includes FREE carry case and 3 year warranty! **SDS5032E**



50MHz SCOPE \$399

50MHz, 4-ch scope at 2-ch price! Up to 1GSa/s rate and huge 12Mpts memory! Innovative "UltraVision" technology for real time wfm recording. FREE carry case! **DS1054Z**



60MHz SCOPE \$349

Best selling 60MHz, 2-ch scope with 500MSa/s rate plus huge 10MSa memory! 8-in color TFT-LCD. Includes FREE carry case and 3 year warranty! **SDS6062V**



70-300MHz SCOPES \$839+

Fast, versatile 2-ch 2GSa/s scopes with 8-in WVGA LCD, integrated generator, 14Mpt memory, very low noise floor. FREE carry case available! **DS2000A series**

FREE TECH SUPPORT, GREAT CUSTOMER SERVICE 

Low Power ARM Module

- Atmel ARM Cortex A5 536MHz Processor

SoM-A5D36

- 4GB of eMMC Flash
- 512 MB of LP DDR2 RAM
- 16MB of Serial Data Flash
- 22 GPIO (3.3V) Lines
- 6x Serial Ports
- 24-bit LCD Controller
- Up to 720P Video
- Touch Controller
- External Address/Data Bus
- Internal Real time clock/calendar
- 4 PWM Channels, 5 Timer/Counters
- 10/100/1000 BaseT Ethernet
- 2x USB 2.0 High Speed Host ports
- 1x USB 2.0 High Speed Host/Device port
- 6 channels of 12 bit A/D (0 to 3.3V)
- 200 pin SODIMM form factor (2.66" x 2.375")



Industrial Temperature



Designed and manufactured in the USA, the SoM-A5D36 is a System on Module (SoM) based on the Atmel ARM Cortex A5 ATSAMA5D36 processor. This low power, wide temperature ARM 536 MHz SoM utilizes 4GB of eMMC Flash, 16MB of serial data flash, and up to 512MB of LP DDR2 RAM. Like other modules in EMAC's SoM product line, the SoM-A5D36 is designed to plug into a custom or off-the-shelf carrier board containing all the connectors and any additional I/O components that may be required. Qty 1 pricing is \$155. Please contact EMAC for OEM & Distributor Pricing.

http://www.emacinc.com/products/system_on_module/SoM-A5D36

Since 1985
OVER
30
YEARS OF
SINGLE BOARD
SOLUTIONS

EMAC, inc.
EQUIPMENT MONITOR AND CONTROL

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

about 110°F (with the heatsink). With a 51- Ω load, the output is 10 V_{pp} and the temperature rises to a reasonable 141°F. Running into 10 Ω provides 3 V_{pp} and a temp of 162°F. A short on the output pushes the temperature up to 170°F. (All of these measures had an ambient temperature around 70°F.) It is not recommended to use a load of less than 50 Ω . Note that a socket is required; otherwise, it may not be possible to slide the heat-sink under the amplifiers.

Since digital pots do not work well at high frequencies (above 1 MHz), manual pots are needed to set the output level. And since these pots are in parallel with the digital pots,




PHOTO 4

Shielded wire and good connectors attach the PCB to the outside world.

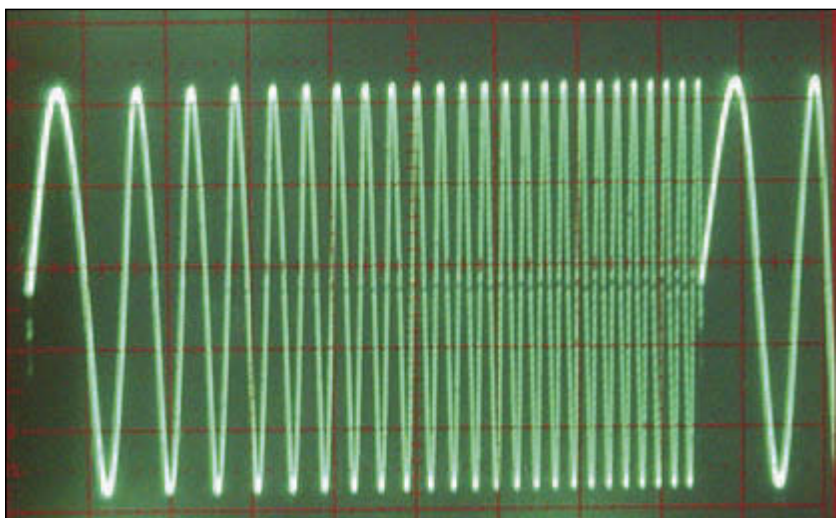


PHOTO 5

A digital sweep looks a lot like an analog sweep.

the setting of one can affect the level of the other. For that reason a switch was added. If you choose to incorporate only one amplifier, the switch can be eliminated and jumpers can be used instead (see **Photo 1**).

Note that the clock oscillators are independent of the amplifier. You can choose to use either the 5- or 50-MHz clock with either amplifier (via software). Or, you can choose to use only one oscillator and not install the other. Also, the datasheet, schematic, and PCB show C6 and C7. These filter out clock noise from the AD5930. However, they also attenuate high-frequency signals (greater than 10 MHz). Since the low-frequency option doesn't work at those frequencies, they can be omitted. And since the high-frequency amplifier has its own filter, they again can be omitted. Positions are left for them on the PCB just in case of some future need.

CONSTRUCTION

First, you need to decide which options you want. If you don't want the high-frequency option, you can eliminate the parts in the box in Figure 1 and use two jumpers for SW1. If you don't want the low-frequency option, you can eliminate U3, U4, and the associated parts. Again, you can use two jumpers for SW1. You can choose to operate with only one clock oscillator, either 5 or 50 MHz. The 5-MHz clock provides 0.3-Hz steps and has a Nyquist limit of 2.5 MHz (usable to about 1 to 1.5 MHz). The 50-MHz clock has 3-Hz steps and a Nyquist limit of 25 MHz (usable to about 10 to 15 MHz).

RXXX is used if only one oscillator is present. This lets the software know which clock is available. It should be 332K to enable the 50 MHz clock and 182K to enable the 5 MHz clock. If it's absent, both oscillators are assumed to be available and are software selectable.

Most of the construction is pretty straightforward but there are several important points to note. The first is that the inline socket-strips that connect to the base-unit are mounted on the bottom of the board (see **Photo 2**). Note that J1 is actually made up of five separate socket strips. The reason for this is space. Using a full row of pins on the inside makes routing difficult and the unused pads take up a surprising amount of space.

The second point is that U1 and U5 are only available as surface-mount parts (SMT). SMT part U5 is the stereo amplifier and is an eight-pin SOIC. It has leads on 0.050" centers which is half the spacing as standard DIP through-hole parts. Most designers won't have too much trouble soldering that. However, U1 the sweep-generator chip, has a 20-pin, TSSOP footprint with the leads on

SOLDERING FINE-PITCH SMT DEVICES

Surface mount is here to stay. And, quite simply, you can either learn to work with these parts or get left behind. There have been similar changes in the past: from tubes to transistors and from transistors to integrated circuits. Of course, there will be more changes in the future. "It's like deja-vu all over again" (Yogi Berra).

Some SMT (Surface Mount Technology) parts are simply impossible for the typical hobbyist to work with. For example, soldering leadless chip carriers and ball-grid devices can only be accomplished with some sort of solder-reflow oven. Clearly this not practical for the vast majority of readers. However, manufacturers generally provide at least one version of their product in a leaded package. And it is certainly possible for hobbyists to incorporate these modern devices without too much expense.

SMT FORMATS

There are many different types of SMT packages. Photo 1 shows several as well as a standard Dual In-line Package (DIP) IC. There are lots more.

Small Outline Integrated Circuit (SOIC) parts have a lead pitch of 0.050" (one pin per 0.050"). This is half the spacing as a standard DIP. This can be soldered with a regular fine-point soldering iron and common 0.030" diameter solder if care is taken. However, it is basically impossible to solder any part with a smaller lead pitch this way. So, don't waste your time, parts and energy trying.

Obviously, the important aspect for us is how close the pins are to each other (lead pitch). The actual number of pins and whether the IC has leads on two or four sides isn't that significant. Generally, "fine pitch" parts space the leads from 0.0256" (0.65 mm) to 0.0197" (0.50 mm).

Note that TSSOP (Thin Shrink Small Outline Package) does NOT define the lead spacing. The 20 pin TSSOP in Photo 1 has a lead pitch of 0.65mm while the 48 pin TSSOP's is 0.50 mm. The Shrink Small Outline Package (SSOP) is usually 0.0250" (0.635 mm), which is one quarter the spacing as the standard DIP. While these differences aren't really important for soldering, they are absolutely critical when laying out a PCB.

REQUIRED EQUIPMENT

A soldering iron with adjustable heat is useful but not absolutely required. However, if you are a serious hobbyist




PHOTO 1:

A common assortment of package types

you either have one or should get one. The other two items you need are a "needle" type of soldering point and some very thin solder. That's it.

The first thing to do is to obtain a true "needle-point" soldering tip. You might be able to buy one for your iron, if they are available. They typically cost a few dollars. (Or, you can make one like I did.) The soldering tip at the left in Photo 2 is the standard "fine" tip that is actually about 0.055" (1.4 mm). The tip to the right of that is one that I filed down to a true needle point (measures about 0.015" or 0.4 mm). This is about the same as the 1/64" tips that are now available for some irons.

The second necessity is to use ultra-thin solder. The solder on the far right is 0.062" (1.6 mm) from the good old days when vacuum tubes ruled. The middle solder is 0.28" (0.7 mm) and is the current standard for general work. The solder at the left is 0.015" (0.4 mm) in diameter. It is possible to find this solder in "hobby packs" for a few dollars rather than one-pound spools for \$40. (Radio Shack has 0.015" solder "with silver" for about \$6. The part number is 64-035.) A small amount will solder a lot of TSSOP devices. Note: I haven't tried the more common 0.020" (0.5 mm) solder.

MAKING IT WORK

The two key concepts are: you want to be able to heat one pin and only one pin at a time and apply solder to only that pin. That's why you have to have a needle point and extra-thin solder.

The first step is to melt some solder onto a corner pin (Photo 3a). Obviously take care not to make any solder bridges to the adjacent pad. You will find that a lower heat setting works well.

The next step is probably the hardest part of the whole procedure. You have to re-heat only the solder and carefully align the chip to the rest of the pads (Photo 3b). For a small PCB, like the one shown, I find it useful to use some double-sided tape to hold the PCB onto something. I use card-stock. But you can tape it to your desk or a piece of cardboard or whatever suits you. Trying to adjust the part with a moving PCB is simply making a difficult task harder. The alignment can be very slightly off if there is enough flexibility in the one fixed pin to allow alignment with the rest of the pins. But, obviously, it's better to be as perfect as possible in this step. It will make all later steps that much easier. Note: it is important to view the registration (alignment) of the IC to the pads from above. If you look at it from an angle it may appear not to be registered properly (see Photo 3c).

Now you can solder the pin that is diagonal to the




PHOTO2:

A comparison of soldering iron tips and solder

previously soldered pin (shown in Photo 3c). You can use whatever method you prefer to do this. However, approach the pad from the open side so there will be less chance of making solder bridges. Note that Photo 3c is seen from a different angle than in Photo 3b. This angle makes it appear that the chip is not registered properly. But as seen from above in Photo 3b, the IC has good alignment with the pads.

Here's the trick to soldering the remaining pins. Place the soldering iron tip vertically onto the small horizontal part of the lead as shown in Photo 3d. This requires a different method of holding the iron and brings the hot metal close your wrist. So be careful.

While keeping the iron vertical and resting on the pad (the weight of the iron is enough), add a little solder (see Photo 3e and Photo 3f). Bring the solder to the pad near the lead. You don't have to touch the lead. The solder will melt on the pad and wick to the lead. Once this happens remove the solder and soldering iron. You don't need, or want, much solder. You can then move on to the next pin to solder.

However, here is another little trick. Don't solder the pins in sequence. This is because the pin you just soldered is still hot. This heat and the proximity of the iron can be hot enough to melt the solder and make a solder bridge. So, skip around a bit. I solder every other pin and then go back to solder the remaining pins. However, if the iron is too hot or if you heat the pin too much you may need to skip two pins or allow the joint to cool before continuing. You'll know things are too hot if you get solder bridges.

DROPPING THE BRIDGE

There will be times when you get solder bridges. It happens. First eliminate as much solder as possible with a copper-braid product (like Solder Wick) from the pads. Often this will take care of the problem. What's more you probably don't need to re-solder these leads. Look closely, most likely you will see a film of solder remaining. If you can't lift the pin with a gentle pry with a razor-knife (like an X-Acto knife) then the solder joint is fine. The amount of solder shown in the photographs is many times what is really needed.

If the bridge remains, then heat one pin as before and draw the knife between the pins. You don't need to press down hard and score the PCB. All you want to do is to separate the solder's surface tension. Once this happens, you will see a space open up. Allow this pin to cool and then heat the other pin. This should draw some of the remaining solder to that pin and open up the space even more. Also, the knife blade should be cold. If it's hot, the solder will just wick around it and not remove the bridge.

PRACTICE HELPS

It doesn't take much to work with fine-pitch devices. You could spend a couple of hundred dollars for a special soldering system. Or you could spend about \$12 and upgrade your existing soldering iron with a new tip and some very thin solder. It does require good close-in work and a steady hand. But a little practice goes a long way.




PHOTO 3

First add some solder to a corner pin (a). Attach the chip (b). Solder the diagonal pin (c). Place the iron tip vertically on the IC lead (d). Add just a little solder to the pad (e). Twenty TSSOP pins with 0.0256" (0.65 mm) lead pitch soldered in less than 10 minutes (f).

0.026" (0.65 mm) centers. This is a quarter of the spacing as DIP parts and ordinary soldering tools and methods simply won't work. (Refer to the sidebar, "Soldering Fine-Pitch SMT Devices," for additional information about SMT soldering.)

To make things easier, I expanded the 20-pin TSSOP footprint to a 20-pin, wide DIP footprint. This allows the use of a small SMT adapter board with SIP socket strips (see **Photo 3**). For prototyping it was necessary to be able to remove the chip as required. But many find it easier to solder to a smaller and inexpensive PCB than a larger and costlier one. Whatever you choose, always mount the SMT parts on the PCB first. Their small footprints make soldering it hard. Don't make impossible by putting obstacles in the way. The rule of thumb that says to mount the parts from low to high (height-wise) is a good rule here. The base unit socket strip connectors are considered "low" because only the ends of the pins stick up on the component side of the PCB.

You have lots of options for wiring to the side panel. I choose to use two, quad RCA panel mount jacks for simplicity and low cost (see **Photo 4**). You can use whatever connectors that suit you best. Shielded wiring is not absolutely necessary, but the performance is a little better at the high frequencies. I chopped up a couple of spare audio cables. It was considerably cheaper than buying a lot of shielded wire. Flexibility of the wire is probably the most important consideration.

You can use SIP socket-strips to mate with the connectors. They are cheap and easy to use. However, they are not very rugged. Plugging them in and removing them a few times caused problems. There's nothing worse than an intermittent connector. I chose to go with quality socket-pins and shells. The cost wasn't very much but it did take considerable time to properly assemble them (see **Photo 4**).

OPERATION

The operation of a DDS sweep generator is different from an analog one. An analog sweeper is basically a voltage-controlled oscillator (VCO) governed by a ramp generator. The start and end voltages of the ramp are adjustable and define the frequency range of the sweep. The speed of the ramp is also adjustable and controls the sweep rate. As you can see, it's pretty straightforward.

The DDS sweep generator is essentially a small computer in itself. It's completely controlled with digital words. In order to create what you want, you have to load eight, 12-bit registers. The details of the registers are available in the Analog Devices AD5930 datasheet, as well as from the Operator's




PHOTO 6

In burst mode, the signal is only on for part of the step.

Manual, which is posted on the *Circuit Cellar* FTP site. But a general overview is useful.

Instead of specifying a start and end point (like an analog sweep generator), the start frequency, step size, and number of steps define the sweep range. Since the start frequency and step size are 24 bits, they require two registers each. So, these five registers precisely define the sweep range.

Note that the sweep is in discrete frequency steps rather than a continuous sweep. If the steps are close together, there isn't much practical difference between an analog sweep and a digital sweep (see **Photo 5**). Of course,



PHOTO 7

As a fixed-frequency generator, it's pretty accurate. This 2-PPM accuracy is much better than the specified 50 PPM.



PHOTO 8

Marker identifies frequency of interest




PHOTO 9

A “sweep” of 10 different frequency segments

you can make the steps as big as you like. For example, you can set your steps precisely on the center of every channel in the AM broadcast band.

The minimum step size is the clock oscillator divided by 2^{24} . For a 50-MHz oscillator, this is 2.98023224 Hz. For a 5-MHz oscillator, the step size is 0.298023224. It is unfortunate that binary and decimal don't mix.

There's another register for the step interval. That is, it specifies how long each step is to be. This defines the overall sweep rate. The number of steps multiplied by the step interval equals the sweep rate. The step interval can be specified in either oscillator clock cycles (linear sweep) or output frequency cycles (log sweep). The step-interval can be as short as two clock cycles (or output cycles) and as long as 1,023,500 cycles. And since you can have as many as 4,095 steps, a full sweep can take as long as 83.8 s (with a 50-MHz clock).

There's a separate register for the burst function. This defines how long the output is active for each step-interval (see **Photo 6**). (Normally, the output is active for the whole step interval.) This can also be either oscillator clock cycles or output frequency cycles (but must agree with the step interval register, above).

The last register is the control register. This controls all the various options and features. A few of the major options are: sine/triangle wave output, sawtooth/triangle sweep, continuous/burst mode, sync-pulse-per-step/sync-pulse-per-sweep, internal/external control, sweep-up at start/sweep-down at start, and others.

The sawtooth/triangle sweep option defines the output profile. A sawtooth sweep begins at the start frequency, sweeps to the end frequency, and then jumps back to the start frequency. A triangle sweep begins at the start frequency, sweeps to the end frequency, and then sweeps back to the start frequency. There is no frequency discontinuity with a triangular sweep profile (however there may be a phase discontinuity).

The internal/external control allows the user to manage aspects of the sweep manually. In particular, the burst length and the stepping of frequencies can be defined with external signals. This allows single-stepping of the frequency sweep which means fixed-frequency operation. With 0.1% distortion and frequency accuracy of 50 parts per million (PPM), it's a pretty good signal generator (see **Photo 7**).

SOFTWARE

The software is a story all by itself. It started out as a simple driver for the sweep generator chip. But over a period of two years



circuitcellar.com/ccmaterials

SOURCES

HDM16216 LCD
Hantronix | www.hantronix.com

Microcontroller
Microchip Technology | www.microchip.com


it became an operating system consisting of about 9,000 lines of code and comments. If you choose to print it out, you'll need over 200 sheets of paper in your printer. It's available for free (non-commercial use) at the *Circuit Cellar* FTP site. The side-panel layout/graphics and comprehensive operator's manual are available, too.

There is one special aspect to the software. It creates a marker/cursor for an oscilloscope. This is extremely useful for determining the precise frequency at a particular point in the sweep (see **Photo 8**). In this case you can quickly plot the impedance curve of a 2.5" speaker and identify the minimum impedance frequency.

Since the marker is software controlled it is completely independent of the sweep generator. Therefore, there are some limitations. First, it can't be set for the first two or last two frequency steps. Second, there is a software latency of about 3 μ s. This means that for proper accuracy, your step length should be longer than this (150 clock cycles at 50 MHz or 15 clock cycles at 5 MHz). Lastly, there is a re-sync period between sweeps of about 15 μ s. This is seen as an inactive period from the end of one sweep to the start of the next sweep. Generally, these limitations are not too significant.

POWERFUL TESTING SOLUTION

This DDS sweep generator is a very flexible and powerful test instrument that is very useful for any type of frequency work. It can do things an analog sweeper can't. For example, **Photo 9** shows a "sweep" where each step is only a fraction of the period. Ten different frequency steps form one cycle. Note the angle change at the end of the cycle/sweep. And the DDS provides digital precision in frequency accuracy as well as low distortion. Once you use a sweeper, you'll wonder how you ever got along without one.

Coming soon: a micro-ohmmeter/micro-voltmeter plug-in board for the base unit. 

ABOUT THE AUTHOR

Gerard Fonte is the principal engineer at The PAK Engineers and has nearly 30 years of hands-on experience that includes missile guidance systems, electronic warfare, gravity navigation, and projects that "don't exist." He has well over 100 publications. Gerard was awarded the Engineers' Council Outstanding Engineering Merit Award in 2006 for his work on Egyptian pyramid construction. He holds a BA degree in Psychology and a MS degree in Natural Science.

Great 16-Bit DAQ Solutions

Easy to Use • Easy to Integrate • Easy to Support

NEW



Ethernet, Multifunction

E-1608

Only \$499

- 8 SE/4 DIFF analog inputs
- 250 kS/s sample rate
- 2 analog outputs
- 8 digital I/O, 1 counter



High-Speed, Multifunction

USB-1608GX-2AO

Only \$799

- 16 SE/8 DIFF analog inputs
- 500 kS/s sample rate
- 2 analog outputs
- 8 digital I/O, 2 counters



Simultaneous Sampling

USB-1608FS-Plus

Only \$399

- 8 simultaneous analog inputs
- 400 kS/s overall sample rate
- 100 kS/s max for any channel
- 8 digital I/O, 1 counter



Low-Cost, Multifunction

USB-231

Only \$249

- 8 SE/4 DIFF analog inputs
- 50 kS/s sample rate
- 2 analog outputs
- 8 digital I/O, 1 counter



High-Channel Count

USB-2533

Only \$1,149

- 64 SE/32 DIFF analog inputs
- 1 MS/s sample rate
- 2 timer outputs
- 24 digital I/O, 4 counters



Simultaneous Sampling

USB-1602HS

Only \$1,999

- 2 SE analog inputs
- 2 MS/s sample rate
- 3 quadrature detectors
- 4 counters

mccdaq.com



Contact us
1.800.234.4232

Infrared Learning

How to Mimic IR Control Signals

FEATURES

Here's a simple way to add infrared control capability to your project. The tiny module can learn and reproduce infrared signals better than most other methods at a fraction of the cost.

By Tommy Tyler (US)



PHOTO 1
Chunghop L102 learning remote

Suppose one of the functions of that wonderful new gadget you're designing is to transmit an infrared (IR) control signal when a certain event occurs. What if you're not sure exactly which signal will be needed or whether it might have to be changed at some future time? I ran across this problem while designing a gadget motivated by an obsessive, almost fanatical dislike of TV commercials. Watching a news program or weather broadcast in real time was always unpleasant during the lengthy, obnoxious, and incredibly repetitious commercials. So my gadget was supposed to turn off TV sound after receiving a UHF signal from a tiny transmitter pinned to my shirt pocket, and then turn it back on after a three-minute commercial break unless I sent the signal again to immediately abort the silence. Nothing rocket science about that, right? To avoid TV modifications or messy wiring connections, I wanted to turn the sound off and on with an infrared mute command learned from the TV remote control. That raised the complexity of my project to a whole new level.

SEARCHING FOR AN EASY WAY

If you're like me you do not relish digging into a problem like this. You just want a simple, reliable solution that doesn't require you to be an expert on IR remote control technology, adds as few parts as necessary to your gadget, and preferably doesn't require programming

a processor. An Internet search turns up a few articles that seem to offer a solution. You might have noticed these solutions usually rely on a three-pin IR receiver module as an optical sensor to learn the IR signal. That sensor combines a fast photo diode, high-gain amplifier, and demodulator in one unit. For maximum sensitivity and rejection of noise from ambient illumination, it also includes a band-pass filter whose center frequency is chosen to match the carrier frequency of the IR signal it must receive. To accommodate various brands of consumer products, receiver modules are manufactured in at least 11 different varieties, with center frequencies of 30.0 to 56.9 kHz. Most IR learners compromise by using a 38-kHz receiver module to detect as many carrier frequencies near that value as possible. They learn a signal by looking at its demodulated waveform without ever actually knowing its carrier frequency. Since the reproduced signal must have a carrier to be usable, they just give it 38 kHz. That might be close enough for some signal protocols like Sony12, NEC1, and RC5 or RC6, but it could be problematic for signals whose carrier frequency is as low as 32 kHz or as high as 57 kHz. Testing to decide how well a one-size-fits-all learning circuit like that meets your needs is almost hopeless. You can try learning a bunch of different IR signals generated with a universal remote control, but unless you have all the equipment those signals control,

there's no way to tell if the learned copies are accurate enough to do the job.

You might have decided from your searches that there is no application-specific IC available for learning and transmitting IR signals. But such a chip does exist! It's the processor used in a Chunghop Electronics Model L102 Learning Remote Control (see **Photo 1**). This remarkable little remote can learn and accurately mimic up to 11 different IR signals of almost any protocol at carrier frequencies from 20 to 80 kHz or more. Learned signals are saved in nonvolatile memory so that nothing is lost when power is off. In this article, I'll explain enough about the Chunghop remote to add reliable and comprehensive IR learning and signaling to your project. You can accomplish this on a low budget without knowing anything about IR remote controls. Another perk is that you'll use less than half a cubic inch of space in your gadget.

Unfortunately, the chip with Chunghop's proprietary OTP firmware installed is not sold separately. But you'll easily find a "Chunghop L102" remote on the Internet for less than \$8 including shipping. The price is a bargain, considering that the remote contains nearly everything you need to solve this problem.

INSIDE A CHUNGHOP REMOTE

Unlike most remote controls that snap together like a bear trap, the L102 is very easy to disassemble for access to its components. Just remove the batteries and six tiny Phillips head screws in back to open the case and drop out the board. Most of the components are grouped together at the lower end of the board (see **Photo 2**). The 16-pin AD009-03 processor IC is the only item that is unique. The eight-pin 24C16 EEPROM and other parts are icing on the cake. If you are designing your own PCB, you can just transplant the processor and EEPROM to your board. But if you're nervous about unsoldering the surface-mount ICs with their tiny, fragile leads, I'll show you how to make a module out of the lower end of the board that provides one intact unit with most of the major components preassembled.

The schematic of this incredibly simple remote shows an outstanding example of well-engineered product design (see **Figure 1**). I've labeled the keys with their generic designations (K1, K2, K3, etc.) rather than Chunghop's arbitrary names (POWER, VOL, CH, etc.). Each key is capable of learning and reproducing any kind of IR signal totally independent of the other 10 keys. Half of the processor's 16 pins (4-11) are dedicated to the keyboard matrix, and three pins (1, 2, and 16) are used for power and ground. That leaves one pin (15) for learning and transmitting IR signals, two pins (13 and 14) for reading and




PHOTO 2
L102 learning remote circuit board

writing to the EEPROM, and two pins (3 and 12) for an LED (D2) that guides you through the learning procedure and shows when the unit is transmitting. The internal clock oscillator requires no external components. IR1 serves as both the IR emitter when transmitting and as a fast photo detector when learning. A key's learned signal is transmitted when the matrix lines that intersect at that key are momentarily bridged. For example, bridging processor pins 9 and 10 would transmit a signal learned for K10.

BUILDING THE IR MODULE

Start by removing both battery contact springs from the board. The Chunghop PCB can be cut easily because it is only 1-mm thick

FIGURE 1
L102 learning remote schematic





PHOTO 3

Lower end of PCB used as module



Phenolic. Scribe across the board several times on both sides with a very sharp knife, about 0.75" from the lower edge (see **Photo 3**). Bend the board at the scribed lines and it should break cleanly. If you feel you need to practice

this, try scribing and breaking somewhere in the upper portion of the board that won't be used. Although it's optional, you can also scribe and break off the ears on both sides so that your module is only about 1.5" long. Smooth the edges with a file or sandpaper.

Attach seven small insulated buss wires A-G to the module for connections to a solderless breadboard, perf board, or PCB (see **Photo 4**). For a minimum footprint, you can mount the module on edge. When used that way, a small dab of epoxy or hot-melt glue on each wire near the edge of the module will stabilize it like a molded SIP component. If you want the module to lie flat, you should transfer electrolytic capacitor C5 from the underside of the board to the top with the other components. Be sure to observe proper lead polarity for C5. Solder wires A and G to the abandoned pads for the battery + and - contact springs. Solder wires B and C to the ends of resistors R6 and R2. Solder wires D, E, and F to the traces of U1 pins 3, 8, and 9.


The wiring diagram for the completed module shows how to connect it to the few additional components used (see **Figure 2**). Holding the setup switch (S) closed initiates learning, the same as with the Chunghop remote. K7 was chosen as the transmit switch simply because the processor pins (8 and 9) for that particular switch are on the end of the IC where access for soldering wires is easiest. Other wiring arrangements using a different switch or several switches can be substituted for the one shown.

Switch S is referenced to ground, so it can be almost any type of switch, a mechanical contact, an open collector or open drain transistor, or a processor output that can be tristated so it floats high. But both sides of the K7 switch (as well as any other keyboard switch used) must be left floating and completely isolated. If you want to switch K7 with an active device instead of a dry circuit contact, a CMOS bilateral analog switch is a logical choice. For a small, single switch in a five-pin SMD package, I suggest an SN74LVC1G66. If you prefer a part in a DIP package, or if you need to use more than one switch, the venerable old CD4066 works well. Don't worry about the resistance of the switches you use. My tests indicate that the Chunghop processor switch inputs will tolerate an ON state resistance of thousands of ohms, no doubt to work with the inefficient conductive rubber switches of the keyboard.

You'll need to supply 3.0 or 3.3 VDC to wires A and G for U1 and U2, but not much current. The module draws less than 50 mA while transmitting and about 1 μ A in Sleep mode. After applying power, it needs a couple of seconds for power-up initialization and signals when it is ready by flashing "dash-

PHOTO 4

Attaching wires to module



dot" once with D2.

TEACHING THE MODULE

There are several different ways you can teach an IR signal to your module, depending on the likelihood it will need to be changed in future. If the chances of a change are negligible, you can enjoy the ultimate in simplicity by omitting the in-circuit learning components (D2, S, and K7) altogether, along with wires C and D to the module. Just learn the needed signal to the Chunghop << button (K7) before you dismantle the remote to make the module. (Be sure to test it first.) It will remain safely stored in the memory chip and ready to use after you've completed the module or transplanted the chip to your gadget.

For strictly personal use, or if there's no convenient place to put the learning components in your gadget, you can use a DIP version of the 24C16 EEPROM. Install an eight-pin socket for it in your gadget and another in a second Chunghop L102 that you save as a learning tool. To reprogram an IR signal, transfer the EEPROM from your gadget to the L102, learn the new signal, and then plug it back into your gadget. Be sure the source and destination power is off before transferring the EEPROM.

But if you are designing a gadget to be used with many different brands of equipment, or perhaps if you just want the security of being able to change the signal easily at any time, the wiring diagram shows everything needed for in-circuit learning. With the "teaching" remote facing IR1, an inch or two away, the learning procedure is as follows:

1. Press switch S until D2 turns on steady, then release.
2. Press button-to-be-learned (on teaching remote) until D2 begins blinking, then release.
3. Press switch K7 until D2 turns on steady again, then release.
4. Repeat steps 2 and 3 for more

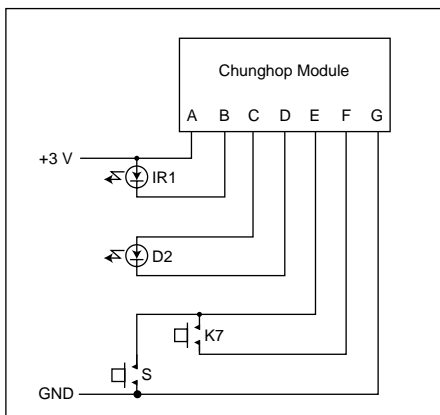


FIGURE 2
Wiring diagram for module

switches, or to correct a mistake.

5. Press switch S to turn off D2 and lock in the learned signal(s).

Many remote controls keep repeating an IR signal as long as a button is held closed. Chunghop will usually transmit just one or two copies of a signal if the teaching sample was very brief and contained only one or two copies. If the teaching sample contained three or more copies, Chunghop will usually mimic that by repeating the signal as long as K7 remains closed. If that causes a problem with your gadget, and you want the module to transmit only one copy of the signal, try tapping the source button very briefly during learning. To test it, hold K7 closed and watch D2. If it gives one short flash it is a non-repeating signal. If it flashes continuously it is treating the learned signal as a repeating signal. Also note that transmission of a signal is never shut off in midstream when K7 is released. The Chunghop processor always completes any signal it has started before turning off.

Although bright red works best, almost any LED of any color can be substituted for D2. The LED you use for IR1 should have an IR wavelength of 940 or 950 nm to work well as a photo detector. For remote transmission up to about 25 feet it's recommended that you salvage and use the Chunghop infrared LED for IR1. If you need

PIC-SERVO

MOTION CONTROL

MOTION CONTROLLERS FOR
BRUSH, BRUSHLESS AND
STEPPER MOTORS.

- controller chips
- controller boards


www.picservo.com

JEFFREY KERR, LLC

FEATURES

EMBEDDED SYSTEM PROGRAMMING

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



A matter-of-fact guide to Assembly that will introduce you to the most fundamental programming language of a processor.

CC-WEBSHOP.COM

ABOUT THE AUTHOR

Tommy Tyler (tomtyler@comcast.net) is a retired engineer with over 40 years of experience in industrial instrumentation, medical electronics, consumer electronics, and robotics product design, earning 17 patents in these fields. He lives in the mountains west of Denver where he pursues his hobbies of technical writing and illustration. Tommy is a contributing expert to the JP1 Forum on IR remote control technology.

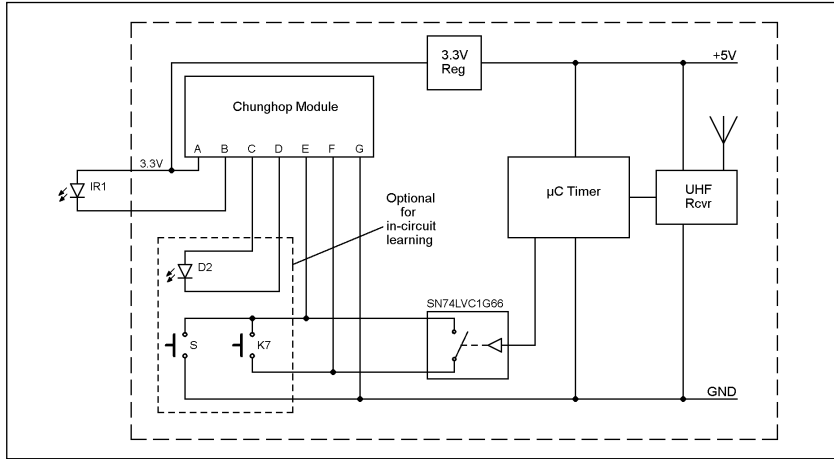


FIGURE 3
The commercial killer


something other than line-of-sight transmission (for example, if the controlled equipment is enclosed in a cabinet, or too far away), place the LED on a long cable like an IR Blaster. Purchased blasters can also be used, but it's hard to predict how well they will perform as a photo detector during learning.

As you can see in **Figure 3**, all the heavy lifting is done by the module. The timer is trivial, and the UHF link uses purchased components. When I convert my breadboard to a printed circuit, I plan to use a 15-position inline socket for the module so it can be easily unplugged for

service or replacement.

A final comment on Chunghop. When searching for the Model L102, you may run across a Model RM-L7, which is a smaller, egg-shaped learning remote with seven buttons instead of 11. It sells for slightly less than the L102, but I don't recommend using it for a number of reasons. The processor is chip-on-board mounted beneath a blob of epoxy, and access to soldering points is much more tedious. It has nonvolatile learning-data memory, but there is no separate EEPROM.

LEARNED SIGNAL ACCURACY

Figure 4 and **Figure 5** show how well Chunghop copies IR signals transmitted by OEM remote controls. This data was captured and recorded with an IR widget. The first line shows measured carrier frequency and a thumbnail picture of the captured signal. The following line(s) show timing details of the demodulated waveform, with carrier ON bursts (high) and OFF times (low). Numbers below the line are time in microseconds, and numbers within the ON bursts are the number of carrier pulses counted during that burst. As you can see, the carrier frequency and all timing of the copy is within a few percent of the original, well within the allowed tolerance of remote controlled equipment. 

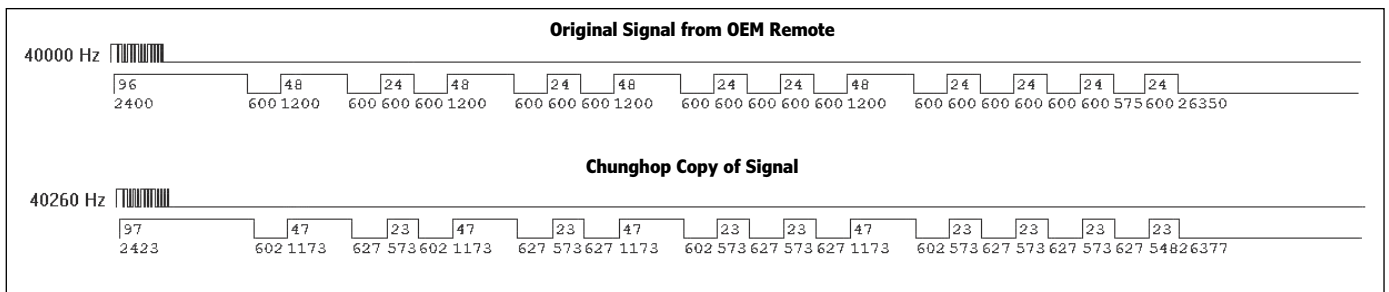


FIGURE 4
Example of learned Sony12 infrared signal with 40-kHz carrier

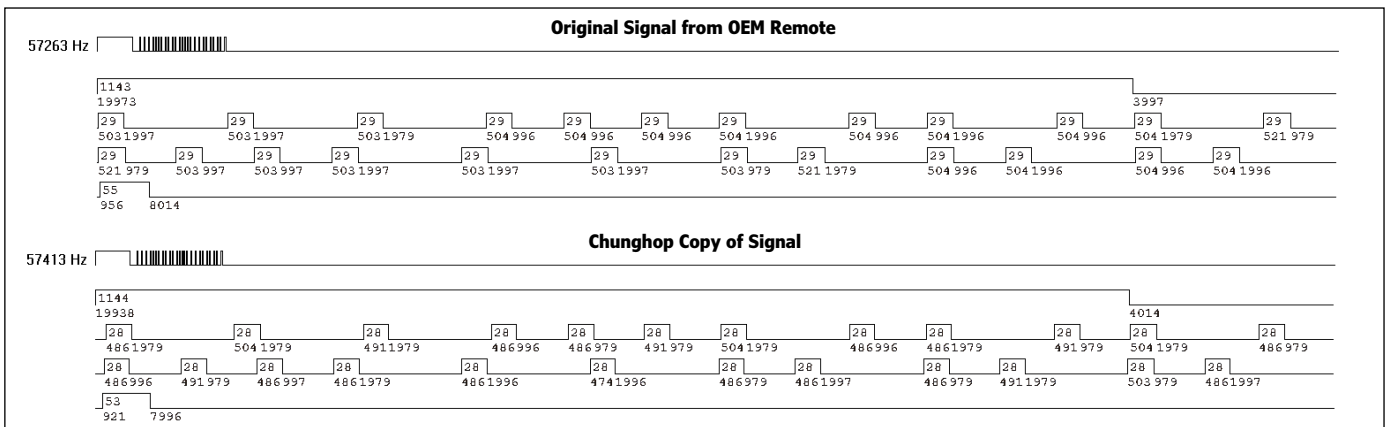



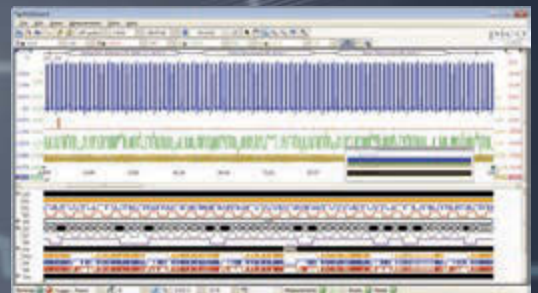
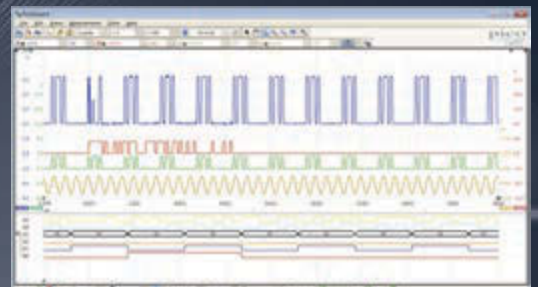
FIGURE 5
Example of learned RCA infrared signal with 57-kHz carrier

MIXED SIGNAL OSCILLOSCOPES

4 ANALOG + 16 DIGITAL CHANNELS

RAPIDLY DEBUG COMPLEX MIXED SIGNAL DESIGNS

- USB 3.0
- ULTRA DEEP MEMORY
- SEGMENTED MEMORY
- RAPID TRIGGERS



INCLUDES AUTOMATIC MEASUREMENTS, SPECTRUM ANALYZER, SDK, ADVANCED TRIGGERS, COLOR PERSISTENCE, SERIAL DECODING (CAN, LIN, RS232, I²C, I²S, FLEXRAY, SPI), MASKS, MATH CHANNELS, ALL AS STANDARD, WITH FREE UPDATES . 5 YEAR WARRANTY

	3204D MSO	3205D MSO	3206D MSO	3404D MSO	3405D MSO	3406D MSO
Channels	2 analog, 16 digital			4 analog, 16 digital		
Bandwidth	60 MHz	100 MHz	200 MHz	60 MHz	100 MHz	200 MHz
Buffer memory	128 MS	256 MS	512 MS	128 MS	256 MS	512 MS
Max. sampling rate	1 GS/s					
Signal generator	Function generator + Arbitrary waveform generator					
Digital inputs	100 MHz max. frequency, 500 MS/s max. sampling rate					

www.picotech.com/pco538

The Sentry Project



PHOTO 1

The Sentry system uses commercial IR motion sensors (lower left) together with a custom vibration sensor (lower right) to determine where an individual is within a building. The base unit (top) integrates reports from these sensors to generate alerts to a caregiver.

The Sentry project uses an array of passive IR sensors placed throughout a building to track motion. The system comprises an RF link to a processor along with an Ethernet module to unobtrusively monitor motion and activity levels.

By David Penrose (US)

My Sentry System is designed to assist those folks living alone who desire the peace of mind provided by a caregiver looking after them without the caregiver having to be present. Its implementation was facilitated by the WIZnet WIZ550io Ethernet module, which provides a rich yet simple interface to the Internet. With a simple microprocessor, the system allows the status of a resident to be continuously monitored in a minimally intrusive fashion. Any abnormal conditions can immediately be alerted to a remote caregiver for action. In this way, a caregiver's smartphone acts as an alert system by letting them know when a resident's activity deviates from a normal

pattern. The system is designed to be simple to set up yet very flexible in its application so the needs of different residents can be addressed. A resident with minimal needs can be monitored by a set of relaxed rules, while a resident in need of more continuous observation can be assigned a set of strict rules. In all cases, the overarching design approach was to provide a system that augments the caregiver's capability.

THE APPROACH

The Sentry System integrates motion sensors, a microprocessor, and the WIZ550io Ethernet interface to monitor a resident and report abnormal activity patterns to a remote

caregiver (see **Photo 1**). The relationship of these subsystems is illustrated in **Figure 1**.

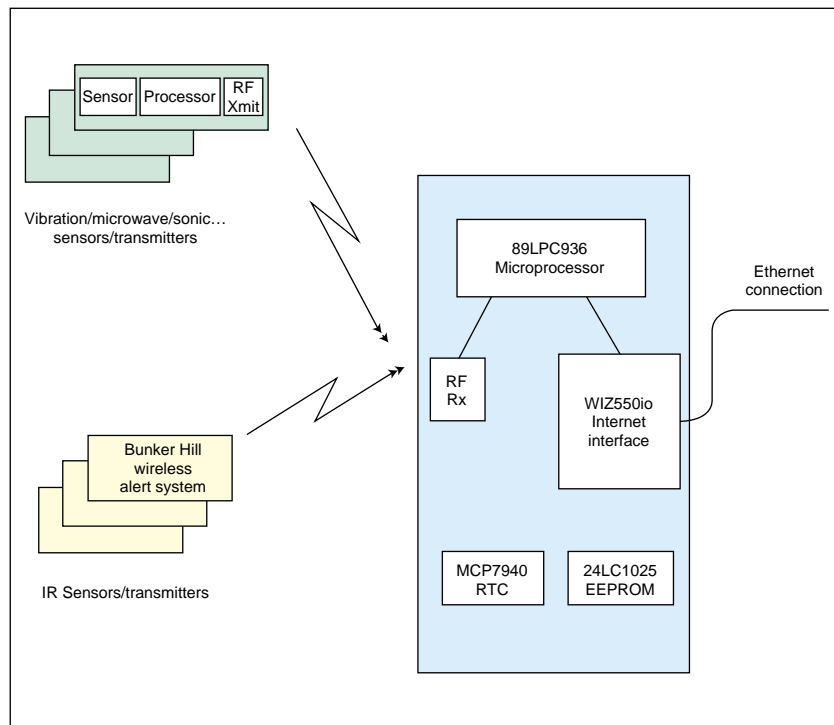
The primary sensors are IR motion sensors. These can be augmented by vibration sensors, pressure mats, ultrasonic, and other devices capable of detecting a person's presence. These sensors are placed at key locations in a resident's home to monitor movement from room to room or within rooms. The vibration sensors are placed in favorite chairs/couches or in the bed to determine if the furniture is occupied and if there is normal activity. All of these sensors are battery powered and report over an RF link. The RF reports from these devices are received by a base unit which then compares the resident's location and activity to a set of rules that define normal behavior for different times of day. Any deviation from normal results in an SMS text message or e-mail being sent to the caregiver along with information about how to contact the resident. In most cases, it is expected that the caregiver would respond by phoning the resident to check on them.

The system is designed to be easy to install and operate. The WIZ550io's Internet interface is used to communicate to a browser allowing the caregiver or resident to configure the system. This configuration consists of identifying sensors and rooms and describing a set of rules for each room for periods in the day. This local interface also allows for a review of all past activity once the system is operational. This history data is valuable for refining the rules to reduce false alarms and ensure security. Since the interface is behind the resident's firewall, the system is secure from improper modification. The key output from the system is the alert to the caregiver, which relies on the WIZ550io module communicating to a service site such as Exosite. The site generates the alerts sent to the caregiver.

SYSTEM HARDWARE

The system's hardware consists of a base unit and multiple sensor/reporting units. The base unit (see **Photo 2**) comprises a WIZ550io Ethernet interface, an inexpensive microprocessor, an RF receiver, a battery backed-up real-time clock, and a serial EEPROM. All of these pieces are integrated into a small form factor case and powered by a plug-in transformer (see **Figure 2**).

The remote units can be one of many different sensor/reporting devices depending on the needs of the resident. The basic sensor is the IR motion sensor, which is available from a number of different sources. I used Bunker Hill Security sensors, which I purchased from Harbor Freight Tools (Item 93068). A sensor plus receiver is very inexpensive.



Some cost only \$11. The item consists of a sensor/transmitter and a receiver/alarm device. The receiver/alarm device is not used in this project although the RF receiver was lifted from one of these units to provide the receiver for the base unit. These sensor units are powered by 9-V batteries and report on an RF link at 433 MHz with a unique address code. The code allows multiple sensors to be deployed and recognized by the base unit.

I constructed a custom sensor from a sensitive vibration-sensing device integrated with a small microprocessor and an RF transmitter (see **Photo 3** and **Figure 3**). This sensor is powered by a 9-V battery and is a working example of just one of the many different types of add-on sensors possible. This particular unit can be attached to a bed or chair and hence continuously report activity even when a resident is not moving

FIGURE 1

Up to eight sensors transmit activity to a base unit processor, which checks for abnormal behavior of a resident. Alerts to a caregiver are generated and communicated over the internet.

PHOTO 2

The base unit incorporates the WIZ550io, an 89LPC936 processor, a MCP79401 real-time clock, and a serial EEPROM to process reports received from the 433-MHz receiver.


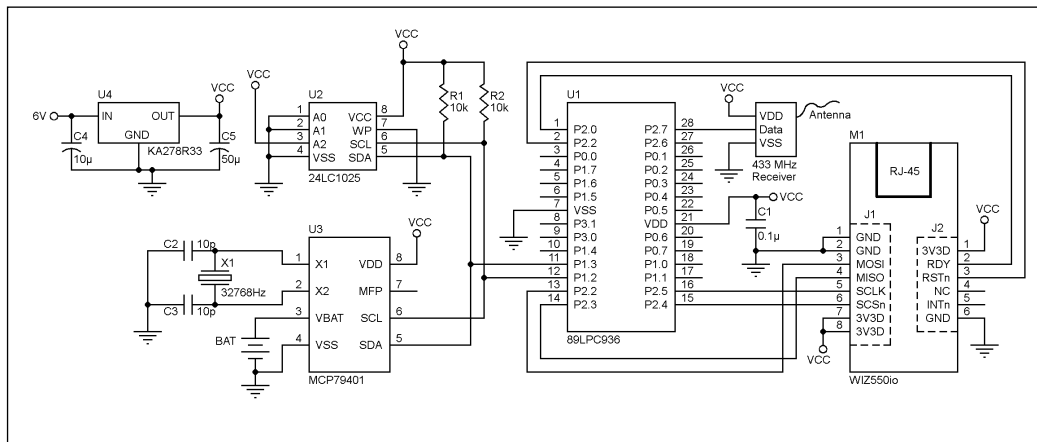


FIGURE 2

The microprocessor accomplishes all of its tasks while using only a few of the available port pins.



about.

BASE UNIT


The base unit integrates the reports received from the remote sensors and processes this data to generate reports issued on the Ethernet interface. This Ethernet interface does all the difficult processing involved in communicating over the Internet so the microprocessor can focus on receiving and processing the sensor data. The processor is assisted by a battery backed up real time clock so that rules for checking on the resident's activity can be

keyed to time of day. An EEPROM stores all the rules and information, including sensor, room, resident, and caregiver information. This EEPROM also stores a time-tagged recording of all status changes and alerts/acknowledgments generated.

The WIZ550io is a key component in the base unit. It features the WIZnet WZ5500 chip and additional hardware to provide a full-featured interface to Ethernet and hence the Internet. The microprocessor talks to this unit over a SPI, while the Ethernet interface is through an RJ-45 socket. There are eight send and eight receive interfaces available within the WIZ550io with a large 16-KB buffer shared among them. This buffer can be allocated in any fashion needed to provide the speed and flexibility required. This large buffer is invaluable for integrating with a small microprocessor since all transmit messages can be constructed in this buffer and all receive information can be read a character at a time from the receive buffer. Without this external memory, using the Internet would be much more difficult or require an external memory to the microprocessor.

The Sentry uses an NXP Semiconductor 89LPC936 microprocessor. It is a very inexpensive but fairly powerful 8-bit processor using the 8051 architecture and instruction set. The 89LPC936 includes 16 KB of flash program memory, 768 bytes of SRAM, 512 bytes of EEPROM, and a wide assortment of peripherals and timers. This project uses the built-in I²C interface, the SPI interface, the capture/compare unit (CCU), and the real-time clock. The CCU in the 89LPC936 provides significant assistance in decoding the incoming RF stream. By just inverting the expected edge with each bit acquisition, the CCU provides very accurate timing of the on/off periods in the data stream.

A Microchip Technology MCP79401 real-time clock chip (RTCC) provides time of day information. The MCP79401 is a Microchip

**PHOTO 3**

This is as simple as it gets: a microprocessor with a vibration switch to wake it up and a 433-MHz transmitter to communicate.

Technology device that implements a battery backed-up real time clock and 64 bytes of battery backed-up RAM. This RAM was valuable for my project, since I needed to keep track of the head and tail of the history file maintained in EEPROM. I did not want to constantly rewrite these values in EEPROM since it has a limited number of write cycles available. I also wanted these pointers to be maintained during power outages so the RAM in the clock chip was the perfect solution. I also use this RAM to store the user's selection of Daylight Saving Time tracking just because it logically belongs with the clock data. I have also used this RAM for a few other key variables but these could have gone equally as well to the EEPROM.

I selected the Microchip Technology 24LC1025 EEPROM for bulk data storage. The 24LC1025 I²C serial EEPROM device provides 1 Mb (128 KB) of storage arranged as 8-bit words in 128-byte pages. I logically divided this chip into two sections of 64 KB each. The lower section stores all the configuration information for my project, while the upper 64 KB is used for a circular buffer of time tagged history data. This chip provides 1 million write cycles with a 3-ms page write time. When evaluating the number of write cycles possible, you need to keep in mind that any write to a page forces the entire page to be rewritten, so don't just write a few bytes consecutively to the chip and expect to achieve the 1 million cycles. One other feature that needs attention is that all writes occur on the 128-byte boundaries. Any attempt to write a block of data across the 128-byte boundary will just wrap around the 128-byte boundary. Individual consecutive writes of single words can avoid this problem but this is slower and the number of write cycles prior to wear-out is decreased. I allocated my use of the EEPROM so all sections fall on a 128-byte boundary and avoided the problem of wrap-around.

Completing the base unit is the receiver subassembly. This 433-MHz receiver is a commodity product available from a number of suppliers. I elected to remove mine from the IR motion sensor system receiver. The device was a good candidate for the receiver in this base unit since it is a self-contained small daughter board in the Harbor Freight units. If you elect not to do this harvesting, these devices are available from a variety of suppliers on the Internet.

THE SENTRY SENSORS

The main sensors for the Sentry system are IR motion sensors. These sensors respond to a warm body moving in their field of view by sending a coded transmission on 433 MHz.

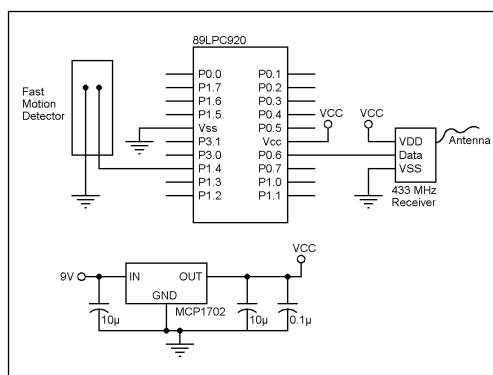


FIGURE 3

The mechanical vibration sensor generates an interrupt to wake the processor, which then formats and transmits a report using the 433-MHz transmitter.

The codes sent are unique to the address of the different units available. This uniqueness of codes allows the Sentry to identify which sensor is reporting and associate the sensor to a room. Setting up the sensor address table is automatic. Each time a new sensor is seen, it is recorded and available for naming and associating to a room.

The on-off keyed sequence of these transmissions is not a very robust data stream and is subject to significant noise. Fortunately, the coding of the transmission sequence allows for noise to be discriminated from the signal by the patterns employed. Only select sequences of 1s and 0s are valid, and the entire sequence must be followed by a 16-ms quiet period. The flipside to this is that a valid signal may be missed by noise interfering with it. To minimize this, the codes are sent a total of 16 times and a full sequence requires about 1.5 s to complete. The base unit processor is designed to ignore repeated transmissions while assembling a valid sequence.

The IR motion sensors can report a resident's location as long as there is significant motion and the resident is in view of the sensor. If the person is sitting down or in bed, then the motion sensors will not be able to discern small movements. To augment the detection capability of the system, vibration monitors have been added which transmit a data stream similar to the IR motion sensors. These sensors are designed to pick up small motions of a healthy individual sitting in a chair or lying in bed. These monitors use a Fast Vibration Sensor (Adafruit 1766) and an 89LPC922 microprocessor transmitting with a WRL-10534 RF transmitter from Sparkfun Electronics.

SOFTWARE

The aforementioned hardware is simple, inexpensive, and readily available. The magic of the project is in its software. The 89LPC936 processor in the base unit provides for 16 KB of code memory in flash memory. The project was developed in assembly language and uses a significant amount of this memory. HTML

ABOUT THE AUTHOR

Dave Penrose (david.penrose@comcast.net) is retired from the aerospace industry. He began his career programming test software for space vehicles and has maintained a life-long passion for developing microprocessor projects. He holds an Expert Class Amateur Radio License (K1DHP), an MS in Computer Science from the University of Santa Clara, and a BA in Mathematics from the University of California at Berkeley. Much credit for putting up with strange devices appearing through the house but also for editing rough thoughts into finished articles goes to Dave's wife, Melinda.

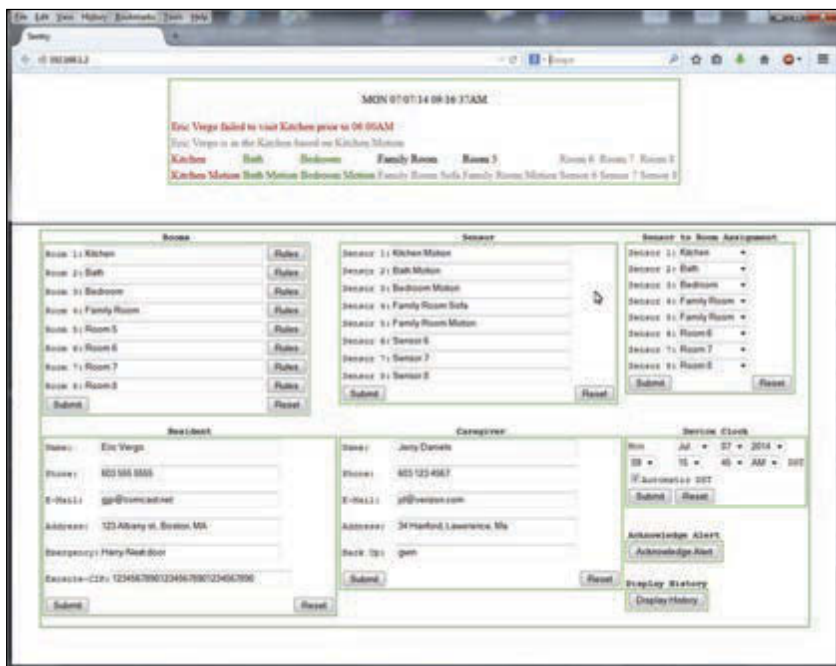


FIGURE 4

The information displayed to a local client allows all aspects of the Sentry system to be defined. It also repeats any alerts that may have been sent to the caregiver.

text strings take a big portion of memory. The software itself uses a main processing loop and two interrupt routines. The main loop begins by initializing the processor, interrupt routines, its peripherals, and key variables that track state conditions. Once the initialization is complete, the project begins its continuous processing loop. At the top of the loop, a 1-s timer is examined. If this flag is set, the MCP79401 RTC is read and used for rule checking and message generation. The next step is to examine a 1-min timer. If this timer flag is set, the logic proceeds to check all occupancy rules that are independent of movement from room to room. These checks will catch a resident being quiet too long or not entering or leaving a room when they should. Any rule violations detected here are flagged for further processing in later steps. The last step of this 1-min tick is to execute

the CLIENT logic, which opens a port to the Exosite system and transmits a message with current time, status, and any alerts.

After the 1-s and 1-min timed events are processed, the loop calls the SERVER routine to listen for any request from a client such as Firefox, Safari, or Chrome. If a request is present, then the appropriate action is taken to either deliver a display page or process data input by the user. This process is called as often as possible to ensure a smooth response to any user requests.

After these events are processed, the main loop then checks for any signals that have been assembled by the interrupt logic from the RF input stream. If data is available, it is identified as to which of the eight sensors it represents, which room it is assigned to, and then the rules are checked for this input. This signal is then removed from the processing stack and the process begins again at the top of the loop.

The signals are assembled for the main processing loop by interrupt logic. This interrupt is generated by the CCU. The CCU in the 89LPC936 is programmed to recognize changes on an input pin and to time durations between edges. Luckily, the CCU also includes a glitch filter so short spikes are eliminated automatically. The job of assembling a sensor report requires that the time duration of the off and on periods are recorded and checked against valid sequences. The interrupt logic first reads the value of the 16-bit timer that has been running in the CCU. It then selects the alternate edge (high to low or low to high) of the interrupt to capture the next interval in the data stream. The time value just recorded is then passed to the DECODE routine to decide if a valid set of data is present.

The DECODE logic first identifies the measured interval as a short low, a short high, a long low, a long high, noise, or a very long interval (15.5 ms), which indicates a sync position in the data stream. The low and high periods are used to assemble a sequence of 1s and 0s, the noise interval resets the process while the sync period indicates that an assembled set can be passed to the main processing loop. The interface between this interrupt logic and the main processing loop is a ring buffer where multiple signals can be stacked up for the main processing loop to use. To reduce the amount of data, this routine will remove repeated transmissions from the same sensor. Each sensor will repeat its message up to 16 times so removing these duplicates early is imperative. This logic will also check for invalid bit sequences in the stream. Once this processing takes place the interrupt will exit until the next edge detected by the CCU.



circuitcellar.com/ccmaterials

89LPC936 Microcontroller
NXP Semiconductors | www.nxp.com

Network Sniffer
Wireshark | www.wireshark.org

WIZ550io Ethernet Module
WIZnet | www.wiznet.co.kr

SOURCES

24LC1025 EEPROM and
MCP79401 RTCC
Microchip Technology |
www.microchip.com

WIZ550IO ETHERNET LOGIC

The WIZ550io accomplishes all of the difficult logic of communicating over the Internet through an Ethernet connection. Without this assistance, the base unit could not possibly accomplish all of its processing and reporting. With this assistance available, the base unit only needs to implement a host to talk to a client such as Firefox or a client to talk to a remote host such as Exosite. Both of these host and client interfaces operate simultaneously in this project and use only a small percentage of the WIZ550io's features.

The host logic in the base unit communicates with a web browser to configure the system. To act as a server, it establishes a Listen configuration and waits for a client browser to establish a connection. The browser will then pass an HTML message requesting a specific page to be displayed or it will pass data to the process from a user form (see **Figure 4**).

The Client processing generates the Exosite remote display and any alert messages to be sent to the caregiver (see **Figure 5**). This process is serviced once per minute and when called establishes a Talk relationship to Exosite. After each interchange of data, this process disconnects from Exosite.

RULE PROCESSING

All rule processing is keyed to rooms in the resident's house. A room can have one or more sensors in it, but all of these sensor reports are combined to determine what room the resident is in and when the last activity in the room took place. There are four simple rules that the caretaker establishes for each room (see **Figure 6**). These rules were based on the logic a person might use to determine if an individual is behaving normally. They can check if a person visits the kitchen or bathroom on a regular basis, if they are in bed prior to a normal time, if they remain in bed past a normal time, if they are too quiet in a room, and other checks that can be made based on time of day and duration in a room. Each rule can be enabled or inhibited depending on what makes sense for a particular room and how the resident is expected to behave.

A violation of a rule generates an alert to the caregiver. This is either an SMS text message or an e-mail saying, for example, "John Doe may need your attention."

NEXT STEPS

The project is complete and working. Prior to commercial use, a feature would be required to use a Domain Name Server (DNS) to retrieve the IP address for Exosite. This is currently hard coded into the program. Secondly, the

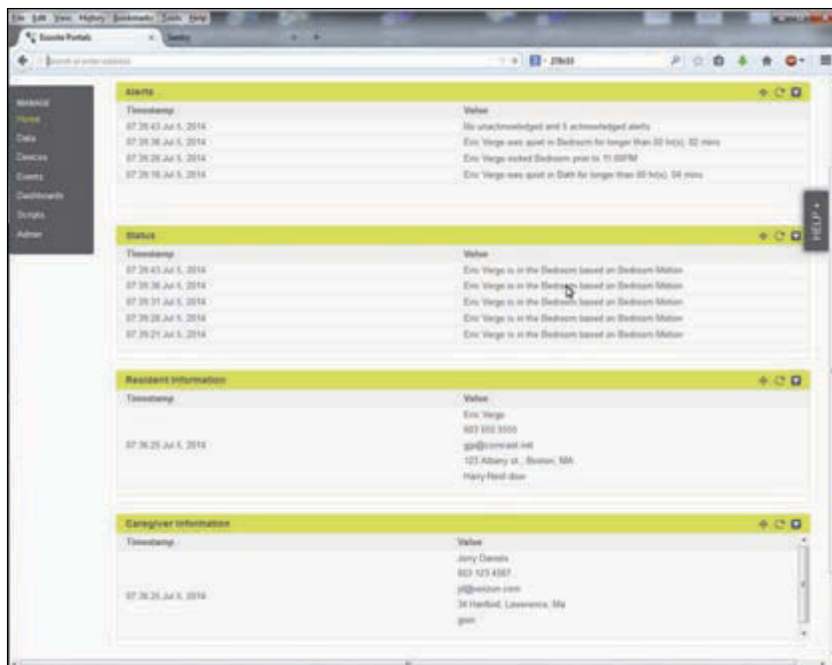



FIGURE 5 Current alert information as well as detailed information on the resident and caregiver is presented on the remote Exosite display.



FIGURE 6 Four simple rules allow the Sentry system to be tailored to any resident's needs.

fixed IP address for the WX550io needs to be fetched from a DHCP server to obtain a local IP address. The first new function I might add would be an audio signal together with an indicator light and push button switch on the base unit. The audio and light would indicate that an alert had been issued. The resident could press the button to acknowledge the alert and let the caregiver know everything is okay. This would be useful in suppressing false alarms but would still record a history of alerts issued.


This project is just one example of how the Internet of Things (IoT) might serve our senior population. Hopefully, it will stir up additional ideas from the talented readers of *Circuit Cellar*. 

Sensorless Maximum Power Point Tracking

FEATURES

Maximum power point trackers (MPPTs) are used to ensure that maximum power is transferred to a device from renewable energy input applications. This article details an impedance-matching method for a closed-loop digital control system, which provides MPPT for input applications that have an approximately fixed internal resistance.

By David Weight, Ian Johnson, Sajjad H. Lalji (UK)



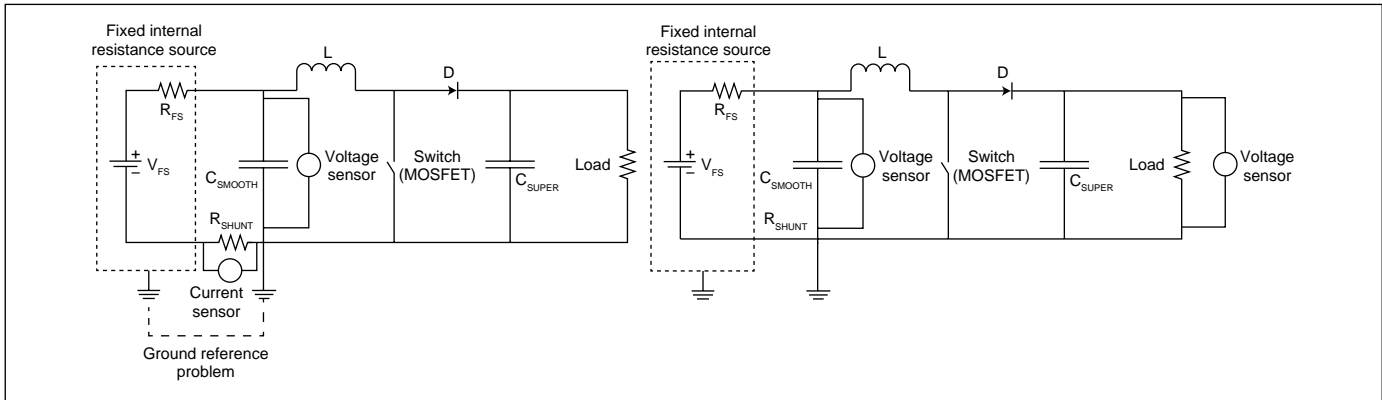
As renewable energy resources become more involved in helping to provide today's increasing demand for energy, energy efficiency is becoming pivotal and thus more importance is being placed on harnessing the maximum power that they are capable of transferring. In order to obtain the highest efficiency possible, maximum power point tracking circuits are used which match the input resistance of the source to the output resistance of the load, a technique known as impedance matching.

The main aim of our project was to develop a new maximum power tracking control method and then subsequently design and construct a prototype converter with a microcontroller used to perform the control. There are a number of existing maximum power tracking methods currently available. However, none of them have approached the problem in the way we do, by using a boost converter operating in discontinuous conduction mode (DCM) with no current sensors used as part of the control algorithm, whilst also allowing for load variation. The trade-off involved in this new approach is that it currently only applies where the source internal impedance is known and approximately fixed. However, once this is established, measuring the input (and output) current becomes unnecessary, therefore removing the need for any current sensors and thereby greatly reducing the cost of the circuit.

We have looked at integrating a new maximum power tracking control method

using a boost converter, without the use of current sensors, which allows for varying input voltage, varying load impedance and varying output voltage when the input power source has an approximately fixed internal resistance (see **Figure 1**). In order to do this, we had to find a new equation to calculate the duty ratio needed to produce the maximum power transfer using only voltage sensors. This was a difficult obstacle to overcome. The duty ratio equations for boost converters found in most textbooks and papers do not look for the maximum power transfer when the variables are constantly changing and involve the voltage and current terms. That's a problem as current isn't measured due to the removal of the current sensor.

Numerous other maximum power point tracking circuits have been designed using boost converters, although many of these incorporate current sensors in their design, with current sensors costing, on average, an additional \$75 or more for a suitable accuracy. Therefore, with the current sensor being the most expensive component used in the circuit, there is a significant cost advantage in removing them from the system. Our circuit cost us around \$75 in total to produce a prototype. The primary renewable-energy applications this can be applied to include thermoelectric generators and inductive power transfer systems, as they have an approximately fixed internal resistance. Unfortunately, the internal impedance of photovoltaic panels varies with environmental



conditions. So this circuit, as it stands, will not function correctly using a photovoltaic panel as the input source. However, thermoelectric generators and inductive power transfer systems are becoming more widely used. For example, if you were to look at one application in particular that would benefit with the addition of this MPPT, it is the new thermoelectric generator being developed by BMW in conjunction with NASA. It improves fuel efficiency and reduces carbon dioxide emissions by up to 5%. In addition, it could eventually replace the need for an alternator, as it aims to produce enough electric power to run all the electrical components of the car. The energy and environmental implications of this will be huge. If successful, it could be applied to all vehicles (cars, lorries, etc.) around the world, which is a massive reduction in oil consumption and a big improvement in energy efficiency as it reuses the heat escaping from the exhaust to power a vehicle. Consider that around 60% of the energy generated by a typical internal combustion engine is lost: half by heat absorbed by the engine cooling system and the other half lost via exhaust heat. Thus, by reusing this exhaust heat, the efficiency of the engine can be significantly improved.

Our prototype circuit features a ChipKIT Max32 development board. It performs the closed-loop digital control that yielded positive results from both simulation and experimental verification for input applications that have an approximately fixed internal resistance, such as thermoelectric generators and inductive power transfer systems.

CIRCUIT OVERVIEW

We decided to use a boost converter operating in the discontinuous conduction mode in order to reduce switching losses. After overcoming the obstacle of creating equations for the duty cycle in terms of voltages only (see Equation 1) and the design equation for inductance (see Equation 2), resulting in the maximum power transfer, we can therefore eliminate current as a required

variable and simply measure the input and output voltages, using these to perform the maximum power tracking.

$$D = \sqrt{\frac{2L(V_o - V_i)}{R_{INT}TV_o}} \quad [1]$$

$$L = \frac{R_{INT}T(V_o - V_i)}{2V_o} \quad [2]$$

This gives us an output of the converter's duty cycle, expressed as a function of the input and output voltage, as well as the component values.

As you can see in Figure 2, the circuit is a very standard generic boost converter. Being a discontinuous converter, it runs in three modes.

In the first mode, the switch is closed and current flows from the source of power, through the inductor, and back to ground. Energy is stored in the inductor during this cycle, and the load is being powered from the energy already present in the capacitor.

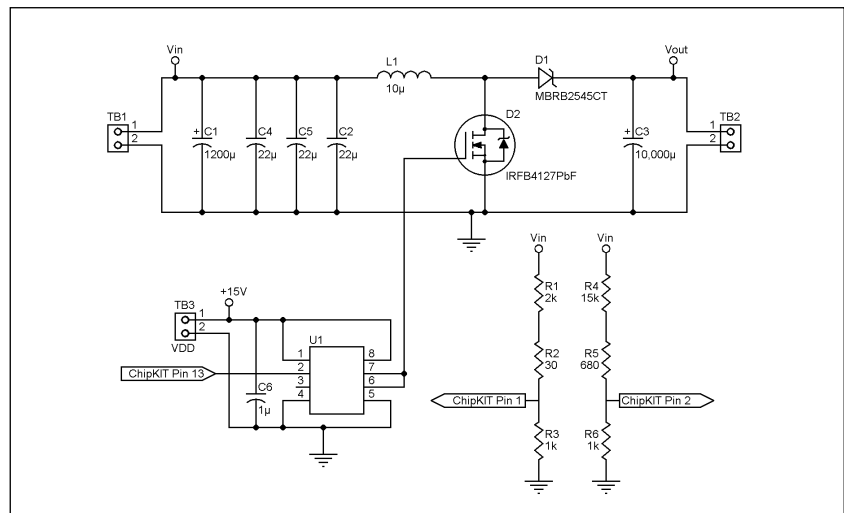
In the second mode, the switch is opened. This suddenly decreases the current flow (or would if it wasn't for the energy stored in the inductor previously). The energy now flows out of the inductor and into the capacitor, which stores energy and smoothes the voltage.

FIGURE 1

A comparison between a traditional MPPT circuit (left) and our circuit (right)

FIGURE 2

The MPPT schematic



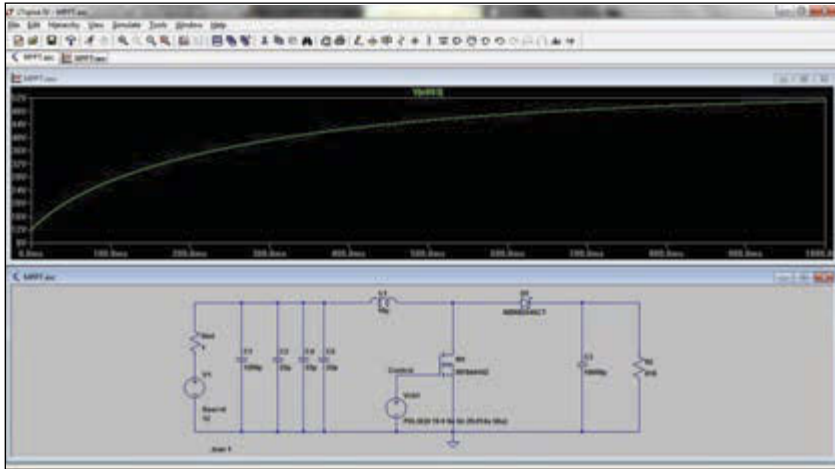


PHOTO 1

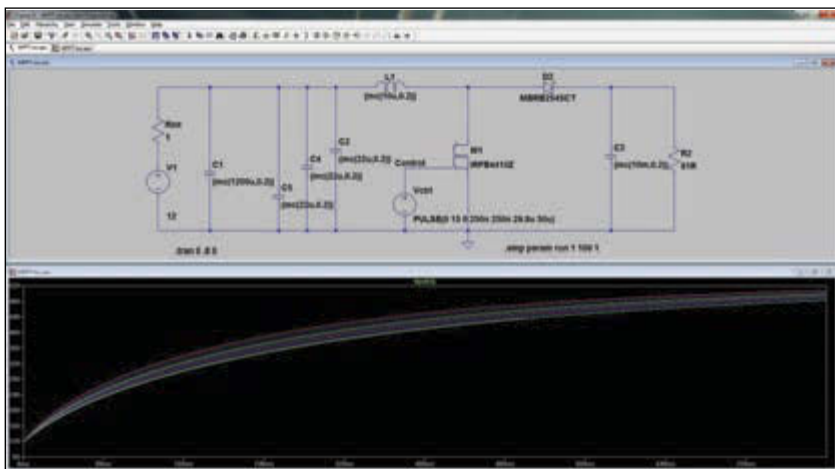
The LTSpice schematic for simulation

ABOUT THE AUTHORS

David Weight (david.weight@wattcircuit.com), Ian Johnson (ian.johnson@wattcircuit.com), and Sajjad H. Lalji (sajjad.lalji@wattcircuit.com) met at Brunel University while studying Electronic and Electrical Engineering as undergraduates. Between their day jobs, they work together on various projects, particularly those relating to LED lighting and embedded systems. They also run the Watt Circuit video blog (<http://wattcircuit.com/>).

PHOTO 2

Monte Carlo analysis of the circuit



In the third mode, the switch remains open and the load runs from the energy stored in the capacitor. The voltage will gradually fall during this phase, until it drops sufficiently such that the circuit enters mode one again, and the whole cycle repeats.

The real control aspect comes from the ChipKIT running the control software. This reads the voltage at the inputs and outputs, and then controls the duty cycle of the MOSFET. As the source internal resistance and the frequency being used are assumed to be known, and the size of the inductor has been chosen based on our inductor design equation, the required duty cycles for different conditions can be precalculated using our duty ration equation and loaded onto the ChipKIT in the form of a look-up table, which makes for extremely fast response using only a simple controller. If dynamic loads need to be considered, it would be a relatively straightforward task to add on board calculation capability to the software.

SIMULATION

Before building circuits, we decided to use SPICE software to simulate the circuit in operation to ensure that the principles seemed

sound. For this, we used Linear Technology's free LTSpice software, an extremely powerful and fast SPICE simulator.

Instead of trying to load the look-up tables into LTSpice and trying to simulate the entire controller, we decided to set the simulation up with a certain input voltage and load, and then set the duty cycle manually. This enabled us to get results extremely quickly (see **Photo 1**).

The results we got from LTSpice looked excellent and gave us a great deal of confidence that the final design should perform as intended. In order to ensure that the design did not look marginal or unstable, we also ran a Monte Carlo analysis on the circuit. This uses the tolerances of the components and varies these on each run, giving an idea of the variation and design robustness of the circuit. The results showed that the circuit seemed to be stable and not overly reliant on any one component tolerance (see **Photo 2**).

CONSTRUCTION

Simulations are all very well, but the next step was to build a prototype and actually take some measurements! We decided to break down the circuit into two parts for the initial prototype. The first section had all the boost circuitry on board (see **Photo 3**). This contains all the circuitry which drives high currents and high-speed switching, and it is therefore pretty dependent on the layout being good and keeping the parasitics minimized. The second part (see **Photo 4**) had the circuitry used to condition the voltage signals in order to make them suitable for the ChipKIT board to read. As this isn't particularly high-speed or sensitive circuitry, we decided to prototype this on veroboard in order to facilitate changes as we refined the design and circuit. Thirdly, the ChipKIT was connected to these two boards and ran the show!

Going through the sections, the first one contains the inductor, capacitors and field effect transistor for the main boost converter. To make sure the transistor performs well, we used a driver IC from Microchip which makes sure that the transistor turns on and off really quickly, meaning it dissipates as little power as possible as heat. This really helps the efficiency of the circuit.

The capacitors are used to filter and smooth the input and output voltages to the circuit. The exact values aren't too critical, but there is a trade-off between cost and how smooth the voltages need to be so we produced a trade-off analysis in order to select the best parts. Having said this, this is a demanding application for the smoothing capacitors as the switching is extremely fast in this circuit, which leads to high frequency components in the current pulses drawn from the supplies.



Slim OLEDs

• Vivid • Organic • Affordable

- Only 5 mm thin
- 2x16, 2x20, and 4x20 character sizes
- Yellow, Blue, Red or Green colors
- Response time: 10 microseconds
- SPI, I²C & Parallel interfaces
- Operating temp: (-40°C ~ +85°C)
- Storage temp: (-40°C ~ +90°C)
- US2066 Built-in controller
- Contrast ratio: 10,000:1
- Viewing angle: >160°
- Low power 2.4 ~ 3.6VDD
- Double height characters
- 3 built-in font tables
- Built-in required logic
- Built-in screen saver



NEWHAVEN DISPLAY INTERNATIONAL

For more information visit: circuitcellar.newhavendisplay.com

Contact us at: **847-844-8795**



RoHS
Compliant





PHOTO 3
The boost converter

The bulk capacitors shown on the top of the PCBs cannot cope with these high frequency components, so we've added several ceramic capacitors in parallel with the bulk capacitors. Although these have much less capacitance than the bulk capacitors, they actually have smaller parasitic components than the bulk capacitors, in particular, equivalent series inductance, and can therefore respond to the high-frequency components.

We used a Schottky, high-current diode as part of the design. Because it is a Schottky diode, it switches reasonably quickly and has a lower forward voltage whilst conducting compared to a standard silicon diode. Typically, a silicon diode will have a forward voltage of approximately 0.6 V, whereas a

PHOTO 4
The voltage scaling circuitry



Schottky diode will conduct at around 0.2 V. This difference in forward voltage translates to gains in efficiency when applied to the entire circuit as this forward voltage translates to energy burned off as heat. In order to provide even higher performance, an improvement would be to replace this diode with another field effect transistor in order to reduce these losses even further. The ChipKIT used for the controller is capable of driving this additional field effect transistor with little added complexity to the software.

The input and output voltages from this circuit need to be fed into the ChipKIT in order to provide the feedback to stabilize the circuit. However, the voltages are too high to feed directly into the ChipKIT as this can only take up to 3.3 V.

We took a simplistic approach for the prototype and used a resistor divider network. This takes the maximum expected voltage and scales it down to the 3.3 V maximum of the ChipKIT. But this has some disadvantages. The resistors have a 1% tolerance and variation here can have a profound effect on the overall performance and efficiency of the circuit. In addition, because this approach wastes much of the dynamic range of the analog to digital converter, there is a loss of precision when the conversion is made.

To address these issues, a couple of improvements are being prototyped. Firstly, as the maximum and minimum expected voltages are known, the analog-to-digital converter does not have to read all the way down to 0 V. Therefore, a differential amplifier can be used to provide an offset to the signal so the entire range of the analog to digital converter is used. In addition to this, a higher-resolution external analog-to-digital converter can be used to replace the on board converter currently used. The on-board device has a resolution of 10 bits, whereas external 16-bit converters are readily available.

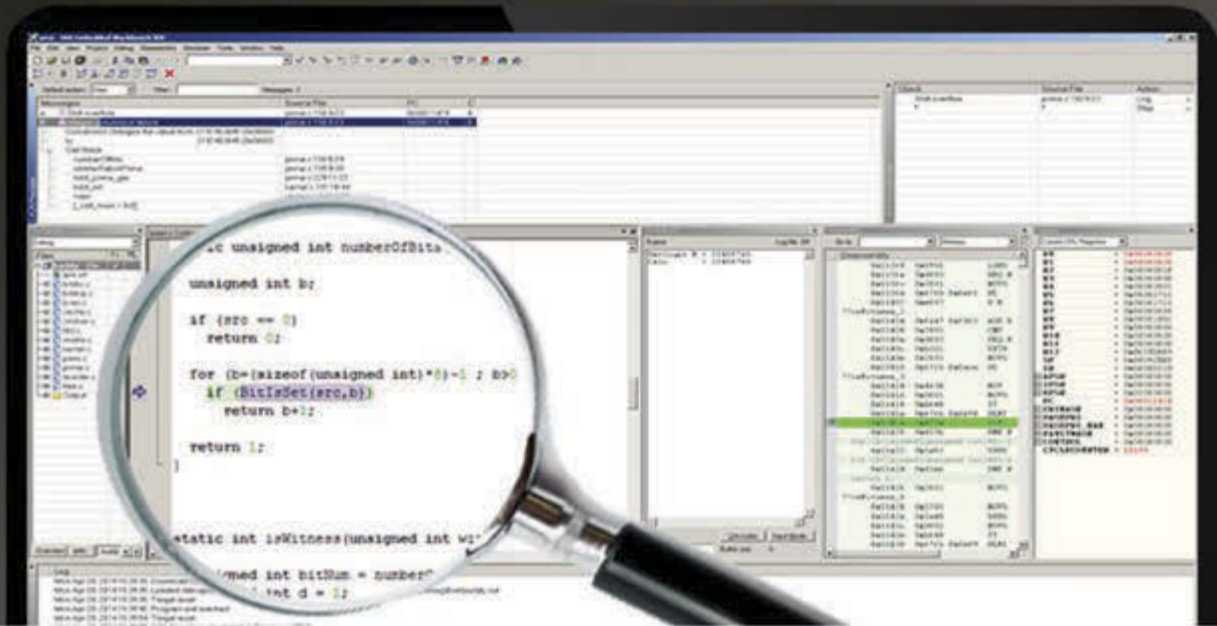
SOFTWARE

The software is really simple. All the hard work was done on a desktop computer with Microsoft Excel to create the look-up table. We used the ChipKIT because it makes software development fast. The ChipKIT MAX32 is a great prototyping platform because of its relatively fast operating frequency and the large amount of memory that is included. This allowed us to add the look-up table into the ChipKIT random access memory for faster operation.

The ChipKIT software is broken up into two main functions, the setup function, and the loop function. Everything inside the setup function is run once upon power-up, here we initialize the input and output pins. At this

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.

www.iar.com/crun





PHOTO 5

The experimental setup

stage we have added the look-up table. We also start a serial connection that is used for debugging and monitoring, this appears as a virtual serial port on your personal computer.

In the loop function, we repeatedly read the voltages at the input pins. Allowing for the fact that these voltages have been scaled down, we then use these as inputs to the lookup table that we defined earlier. Finding the value at this point in the lookup table will be a very quick operation, and this will allow the unit to cope with rapidly changing loads to maintain maximum power point tracking.

The program works by reading the input and output voltages into the ChipKIT through the analog input and output pins. The relevant duty ratio value from the lookup table is returned to perform the maximum power point tracking and send the relevant pulse width modulation signal to the circuit. This

process is then repeated continuously in real time to ensure that the source is operating to provide the maximum power to the load.

TESTING

In order to test the board, we needed to provide a repeatable test setup and it was rapidly decided that using a thermoelectric generator or similar for bench testing would make repeatable testing more difficult. The decision was therefore taken to use a bench power supply unit but to add a precision high-power resistor in order to replicate the source impedance. This resistor is included on our main boost converter printed circuit board just for prototype purposes.

In order to verify the operation of the circuit, we also needed to measure the voltages independently, as well as measuring the input and output current to verify that the circuit is functioning as a maximum power point tracker. Due to budget limitations, we had available Fluke current clamps with a frequency response of only 100 kHz. Unfortunately, this means that many of the oscilloscope captures where current has been captured lacks detail and the switching cycles cannot be seen. The experimental setup is shown in **Photo 5**.

When we tested the circuit, the performance matched the simulations extremely well and surprisingly (for us), the circuit worked first time. After recovering from the surprise, we then tweaked a few resistor values to bring the voltage measurements back in line so they were spot on accurate, which further underlined the need to improve this area of the circuit.

During this phase of testing the debug output from the ChipKIT was extremely useful in verifying the operation of the circuit to ensure that it was in the correct mode and reading the input and output voltages correctly. The measured input and output voltages, as well as the duty cycle selected based on these was displayed on a serial monitor program running on a laptop on the bench.

RESULTS

As you can see in **Photo 6a**, maximum power transfer has indeed been achieved due to the fact that the input voltage to the converter (V_i) is always approximately half of the source voltage (V_S) applied throughout the range of values tested from 2 to 12 V. A more detailed look at the circuit working with a 12-V input voltage and providing an 81- Ω load produced an output voltage of approximately 6 V and can thus be seen in Photo 6b. This was also confirmed over a range of resistive loads ranging from 4 to 400 Ω and an inductive load consisting of a brushed direct current motor. This means that the internal resistance of the source (R_{INT}) always matches the resistance of the converter and thus impedance matching is being performed



circuitcellar.com/ccmaterials

RESOURCES

C. P. Basso, *Switch-Mode Power Supplies: Spice Simulations and Practical Designs*, McGraw-Hill Professional, 2008.

Green Car Congress, "BMW provides an update on waste heat recovery projects; Turbosteamer and the Thermoelectric Generator," 2011, www.greencarcongress.com/2011/08/bmwthermal-20110830.html

I. Laird, et al., "Comparative Study of Maximum Power Point Tracking Algorithms for Thermoelectric Generators," Australasian Universities Power Engineering Conference, 2008.

M. P. Theodoridis and I. Johnson, "Current-sensorless Boost Converter for Maximum Power Tracking of Thermoelectric Generators," Universities Power Engineering Conference (UPEC) 2012, 47th International, 2012.

SOURCES

LTSpice

Linear Technology | www.linear.com

ChipKIT Platform

Microchip Technology & Digilent | www.digilentinc.com

throughout. Therefore, the equations for the duty ratio shown in Equation 1 and for the inductance design value in Equation 2 hold true, as maximum power transfer has been achieved for a full range of values.

In conclusion we have managed to deliver a method for the design and construction of a prototype converter with a microcontroller used to perform maximum power point tracking control of renewable energy devices with an approximately fixed internal resistance, which has been fully analyzed and experimentally verified.

As the source voltage (V_S) is varied, the output voltage (V_O) changes as a result of the system performing maximum power point tracking. The input voltage (V_I) to the converter also changes as it follows half the value of the source voltage (V_S) due to impedance matching (maximum power point tracking). The change in the duty ratio is not static as it is in fact dynamic, which means that the duty ratio is constantly changing until the circuit reaches steady-state, after which point it is restored to its original value. The steady-state value of the duty ratio only changes if the load resistance changes or if the source voltage changes while the output voltage is “clamped” (for instance, if the output was connected to a battery).

Furthermore, this was all made possible thanks to the implementation of digital control, without which specific experiments could not have been carried out. For example, Photo 6a could not have been produced without the automatic updating of the duty ratio from the closed loop digital control system. If you look again at Photo 6a, you can verify that the equation discovered for the duty ratio (Equation 1) and the design equation for the maximum value of inductance (Equation 2) which ensured that the boost converter remained in discontinuous conduction mode indeed holds true. This is done due to the fact that the input voltage is always approximately half of the source voltage when varied over a period of a time. This proves that maximum power transfer has been transferred.

FURTHER WORK

The circuit that has been produced here in order to test this new method of maximum power point tracking can be vastly improved for efficiency and to broaden its applicability. Component selection has only been carried out with a view to circuit prototyping, and therefore large and easily reworkable components have been selected. The performance of the circuit can be easily improved with more modern components. For example, you could replace the Schottky diode with a switched field effect transistor, reducing voltage losses.

Also, the potential dividers used to scale the input and output voltages to the range suitable for the ChipKIT are extremely crude and don't use the full dynamic range of the ADCs. Accuracy, and therefore efficiency, can be improved with a more sophisticated method to scale these voltages (e.g., using differential amplifier circuits instead).

With further modification to our design, we could apply the circuit to a wider range of renewable input sources which have variable internal impedance, particularly if the internal impedance is proportional to other factors. The obvious use is with photovoltaic panels, where the internal impedance can frequently be approximated using additional measurements. By utilizing an irradiance or temperature

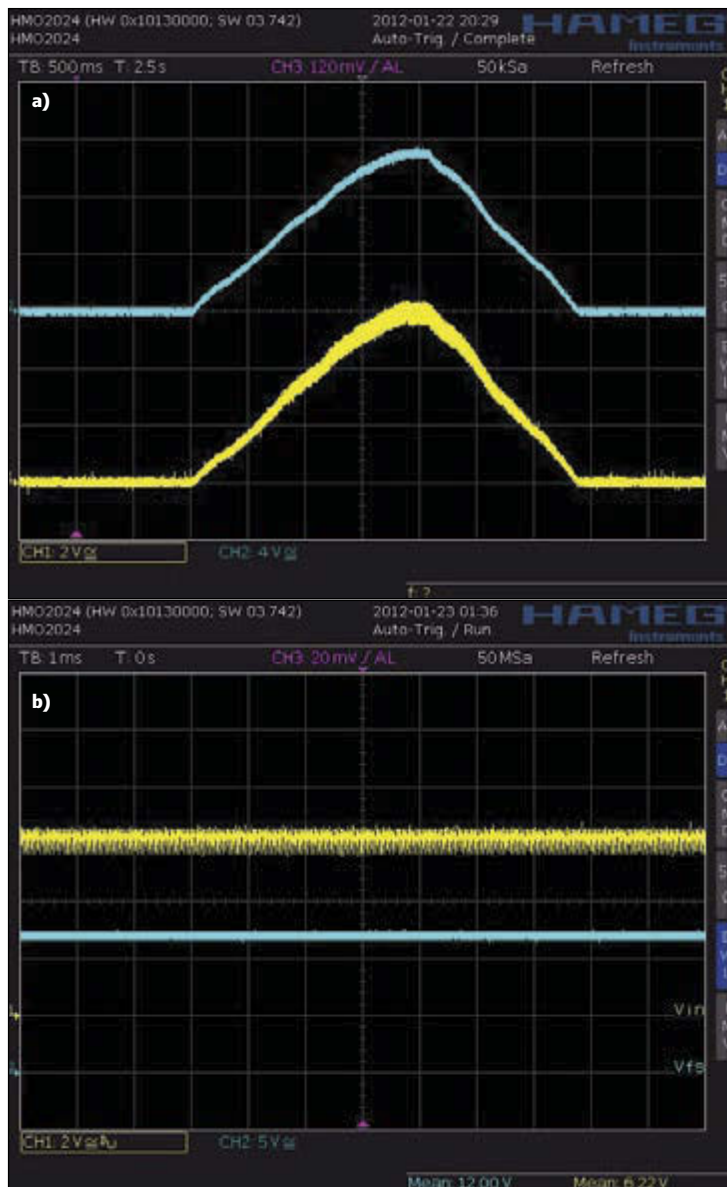



PHOTO 6

This shows circuit maintaining maximum power point tracking as the input voltage is varied (a). The blue trace shows the input voltage and the yellow trace shows the input voltage after the simulated internal resistance. The circuit working with a 12-V input and a 81- Ω load (b). The blue trace shows the input voltage and the yellow trace shows the input voltage after the simulated internal resistance.

sensor combined with the input voltage, the internal impedance can be approximated when the necessary work has been carried out to characterize a particular panel's behavior. This is likely to have poorer performance than a true closed-loop system using a current sensor. But in many applications, the cost/performance trade-off may be attractive as the temperature and irradiance sensors will be cheaper than current sensors for many applications. The variable input and output voltages also favor the use of supercapacitors, which can be used in place of the output capacitor in order to store more energy and charge/discharge faster, in order for the load to be supplied even during times when the input source is unable to provide enough power. 

GREEN COMPUTING

Cooling Awareness in HPC Job Allocation

COLUMNS

High performance computing (HPC) clusters run a variety of computationally intensive applications, each of which typically occupy many processors in parallel. Job allocation in HPC so far has primarily focused on performance. This article discusses how cooling awareness in job allocation can improve energy efficiency substantially.

By Ayse K. Coskun (US)

High-performance computing (HPC) clusters serve large supercomputing applications or, more broadly, applications that are computationally intensive and highly parallelized. Scientific computing applications are common examples of HPC applications. Running highly parallelized applications implies that a job's various tasks run on a large set of processor cores, and these tasks communicate and synchronize through protocols, such as the message passing interface (MPI).

Leveraging MPI-based infrastructures for inter-task communication for most HPC applications indicates that a non-negligible amount of the overall application execution time is spent in sending and receiving messages. In fact, it is common to have 20% to 30% of the overall time spent executing an application to be dedicated to communication.

As a result, current job allocators in HPC clusters seek to group various tasks of a job close together in nearby processor cores. This means that the allocator would prefer cores that are on the same chip or within the same server, and also servers that are within a small number of network "hops" of each other. In this way, sending/receiving messages over many hops, which increases the delay of communication, can be avoided to some extent.

Grouping the tasks of each job close together also helps reduce interference across jobs. In fact, researchers at Sandia

Labs showed a 2x slowdown for two communication-intensive jobs when hand-placing them such that their communication paths overlap significantly (as opposed to grouping tasks of each job close together).^[1]

Optimizing the performance or communication delay of jobs in HPC clusters is a generally well-studied area, and there are many interesting algorithms targeting different types of jobs, communication patterns, or network topologies. A common property of most of these algorithms is to maintain a total "hop-distance" (or total Manhattan-distance) of each job's allocation and then to seek minimizing this distance during job allocation.

Reducing the communication distance of a job helps reduce the time spent in communication while running that job. Thus, one would expect performance-aware algorithms to reduce energy consumption (energy is running time multiplied by average power). One important aspect of an HPC cluster's energy consumption, however, is cooling: 30% to 50% of the total energy is consumed by the cooling infrastructure in current clusters. This number is dropping with more efficient cooling technologies, but is still substantially high.

The rest of this article discusses how bringing cooling awareness into HPC job allocation can provide significant reduction in cooling costs and improvement in overall energy efficiency.

THERMAL DYNAMICS

Figure 1 shows the layout of a typical computing cluster. In this arrangement, rack inlets where the cool air is supplied are facing the outer aisles forming cold aisles at the sides. Rack outlets, where the hot air exits, are facing each other forming a hot aisle in between the two rows. This arrangement is frequently used in modern clusters and is referred to as “hot aisle/cold aisle” layout.

The significance of a hot/cold aisle layout is that some of the exhaust heat gets recirculated back into the cluster. In other words, the hot air outcome of the servers is not fully absorbed by the computer room air conditioners (CRACs). Because of this recirculation and the inherent asymmetries of the layout (e.g., CRAC units residing in certain specific locations of the cluster), even when the servers are all of the same kind, one can speak of a thermal heterogeneity across the servers.

Figure 2 shows the cross-interference coefficient matrix for an example 40-node cluster (where each node may include multiple servers) in a 3-D map. This matrix represents the recirculation phenomena by quantifying the fraction of output heat from each node that is recirculated to the inlet of the other nodes. It is an $n \times n$ matrix for a cluster with n nodes. This recirculation matrix can be obtained for a given cluster through detailed computational fluid dynamics (CFD) simulators or through measurements.

What does such a matrix tell us? Simply, nodes are not equal in their thermal efficiency. For example, some of the compute nodes may be more prone to getting extra heat input because of their relative location to other nodes and to CRAC units in the cluster. As a result of the recirculation and, also, considering there are often idle nodes even in busiest computing clusters, which nodes we select for allocating a job makes a difference in the overall temperature of the cluster.

The recirculation matrix and the power consumption of each node determine the temperature for each server. Server temperatures and the CRAC unit properties then determine how much cooling power is needed to cool down a given cluster. So, job allocation has an impact on the cooling power consumed in a cluster. (Temperature and cooling power dynamics of clusters can be estimated via detailed CFD tools or simpler mathematical models. Some examples of estimation mechanisms are given in recent work.^[2,3]).

EFFICIENT JOB ALLOCATION

How can we make the job allocation process reduce its potential adverse effects on the




FIGURE 1


Layout of a two-row cluster with hot aisle/cold design. Some of the heat exhaust is recirculated back into the system, which creates thermal heterogeneity among the compute nodes.

cooling cost? One straightforward approach would be to prioritize more “thermally-efficient” nodes while selecting the nodes for a job that is waiting to be allocated. This is a complex problem as the solution space is large and also each job allocation decision affects the later decisions. One way to solve this problem is to set up a formal optimization framework with goals and constraints.^[3]

In my research lab, my students and I noticed that while methods to optimize the cooling energy and the performance (communication) cost separately exist in the literature, jointly targeting these goals has not been addressed. If we accomplish to address both goals, this joint optimization can bring energy efficiency benefits. These two goals of reducing cooling and reducing the hop-distance of each job may conflict depending on the specific case. Thus, we decided to set up an optimizer that weighs both goals and

FIGURE 2

An example cross-interference matrix for a 40-node cluster. The 3-D map demonstrates how much each node impacts other nodes in terms of the recirculated heat. A higher bar implies a higher amount of heat being recirculated from a node’s output into another’s inlet.



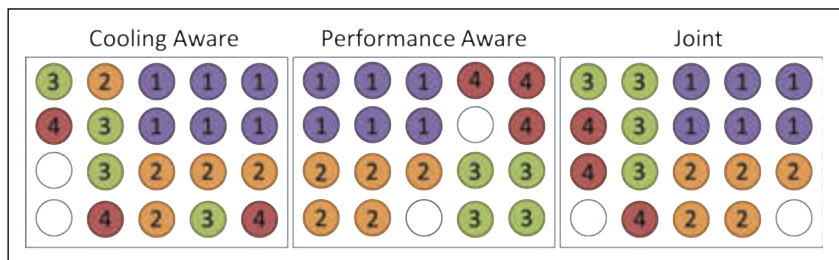


FIGURE 3

Joint optimization policy simultaneously seeks to reduce the cooling cost and also to reduce the communication distance of each job.

minimizes a weighted sum of performance degradation and cooling cost.

Figure 3 shows a simple example to demonstrate the principles of solely cooling-aware, solely performance-aware, and “joint” optimization while allocating four jobs sequentially onto a cluster row including 20 nodes. Each job requires six, five, four, and three nodes to run on. Colored circles represent the busy nodes while the white

circles imply free nodes. The numbers indicate the arrival order of the jobs. For example, job 1 arrives first and runs on six nodes. Cooling-aware policy assigns the jobs to nodes while avoiding the nodes that contribute more to heat recirculation as much as possible. (Note that the sequence here is specific to the example cluster and will change for a different cluster.) On the other hand, performance-aware policy starts from an arbitrary node of the data center and tries to confine the allocated nodes for each job to the smallest total hop-distance possible. The joint policy prefers the cooling-efficient nodes first but also assigns the tasks of a job in close proximity to each other.

In my research lab at Boston University, we expanded this idea of jointly optimizing for performance and cooling energy, and put together a full-scale optimizer that is able to evaluate the performance and cooling costs for an arbitrary cluster and job sequence. This optimization problem, however, may take a lot of time to solve for a large number of nodes because performance and cooling cost problems are both time consuming. For example, a typical HPC cluster has hundreds or thousands of nodes, which creates a huge possible solution space that needs to be spanned.

IMPROVING SCALABILITY & PERFORMANCE

Several methods can be applied for designing an algorithm that can determine the job allocation to reduce both communication time and cooling energy in a reasonable time and complexity. First, one can design heuristics that follow a given set of rules of thumb instead of finding the optimal solution.^[2] Second, it is possible to use lower cost estimators for the thermal evaluations (i.e., linearized models instead of complex CFD simulations).

In a recent on-going project, my students and I formulated the joint optimization problem as a binary quadratic programming problem, which is actually a combinatorial optimization problem. It is an NP-hard problem; however, in practice, it can be efficiently solved using well-known discrete optimization techniques such as the branch and bound algorithm. For example, it is possible to solve this optimization problem using the TOMLAB/CPLEX solver using existing functionality in the tool.

Figure 4a shows the average increase in job execution time for each policy in comparison to ideal job running times (ideal case considers optimal communication delay for each job). Even the performance-aware policy results in a slight increase in job running times. This is because the job running times in the ideal case do not consider the utilization of the cluster

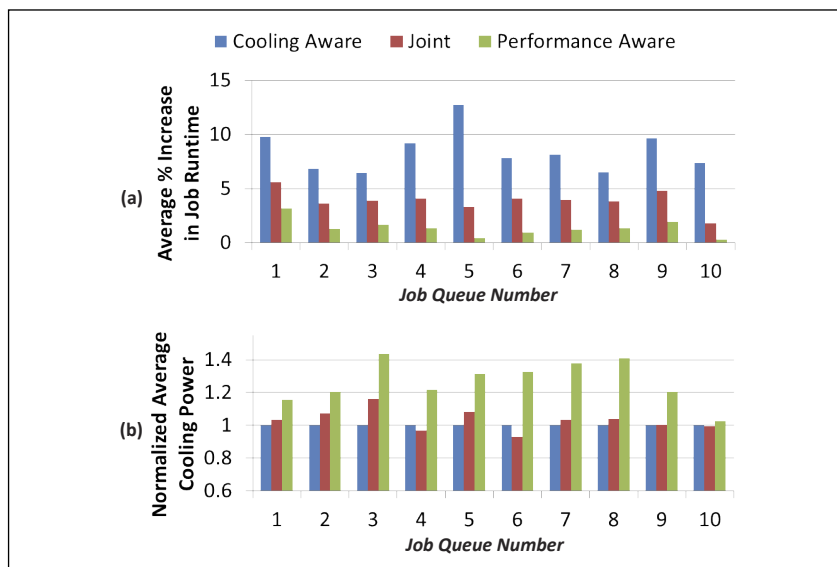


FIGURE 4

Comparison of running times (i.e., larger running time means lower performance) and average cooling power for various job queues arriving at a 40-node cluster including 800 processors. The average communication time to running time ratio for each job (before applying any allocation) is 30% in this example.



circuitcellar.com/ccmaterials

IEEE International Conference on Distributed Computing Systems, 2002.

[2] F. Kaplan, J. Meng, and A. K. Coskun, “Optimizing Communication and Cooling Costs in HPC Data Centers via Intelligent Job Allocation,” Proceedings of International Green Computing Conference (IGCC), 2013.

[3] Q. Tang, S. K. S. Gupta, and G. Varsamopoulos, “Energy-Efficient Thermal-Aware Task Scheduling for Homogeneous High-Performance Computing Data Centers: A Cyber-Physical Approach,” IEEE Transactions on Parallel Distributed Systems 19, 11, 2008.

REFERENCES

[1] V. Leung, et al., “Processor Allocation on Cplant: Achieving General Processor Locality Using One-Dimensional Allocation Strategies,”

and the availability of the nodes. In other words, depending on the current utilization of the data center and the locations of the available nodes, performance-aware policy may not always guarantee an ideal allocation for each job. Cooling-aware policy, however, results in considerably longer job running times, resulting in an up to a 12% increase for the given job queues. Intuitively, jobs with a lower communication to computation ratio will see a smaller increase in their execution times; however, the general trends shown here would still be valid.

Figure 4b compares the average cooling power for the three policies across the same job queues as Figure 4a. The joint optimization policy results in an average cooling power close to the power consumption achieved by the solely cooling-aware policy. In comparison to the solely performance-aware policy, joint allocation saves 18% cooling power on average. An interesting observation is that for some cases, the joint policy achieves lower average cooling power than the cooling-aware policy such as in queues 4 and 6. This is because with the cooling-aware allocation, some jobs may take a longer time to run in comparison to the joint policy (as the allocation is not performance-optimized). When a large job is running for an extended period of time, it keeps the nodes active for a longer time. This results in higher total computing power and cooling energy in comparison to the joint policy.

Another significant aspect of building job allocation optimization policies is the access to simulators that can evaluate realistic scenarios. This is especially critical as many researchers do not have privileges to test job allocation policies on real-life clusters. My students and I were initially using our in-house simulators, but we are now expanding the open-source SST simulator (<http://sst.sandia.gov/>) with power and cooling models. The SST simulator is able to run real-life job traces collected from various machines under a set of different job allocators/network topologies.

FUTURE TOPICS


Many open problems remain to be solved in the area of efficient HPC job allocation. Even though future data centers are likely to be more cooling efficient (via the use of in-row coolers, more substantial use of liquid cooling, etc.), which means less recirculation effects, the servers in cluster are still going to be thermally heterogeneous. This is because a cluster often includes a variety of hardware, and also, asymmetries in cooling still exist in many recently built clusters. Thus, taking thermal efficiency into account during job allocation can help push the performance of

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. Coskun'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. Coskun serves as an associate editor of the *IEEE Embedded Systems Letters*.



the HPC applications.

Other open problems include combining job allocation with power budgeting. (I talked about power budgeting in my *Circuit Cellar* 292 article, "Budgeting Power in Data Centers.") In my lab, we are also investigating how to combine job allocation with task mapping, which is deciding where to allocate each specific task of a given job. Task mapping can bring additional performance and efficiency benefits especially for jobs with largely diverse tasks. 

Get PUBLISHED. Get NOTICED. Get PAID.


Circuit Cellar feature articles are contributed by professional engineers, academics, and students from around the globe.

Each month, the editorial staff reviews dozens of article proposals and submissions. Only the best make it into the pages of this internationally respected magazine.

Do you have what it takes?

Contact **C. J. Abate, Editor-in-Chief**, today to discuss the embedded design projects and programming applications you've been working on and your article could be featured in an upcoming issue or online at circuitcellar.com.

Email: editor@circuitcellar.com



THE CONSUMMATE ENGINEER

Essential Electromagnetic Compliance (Part 4)

COLUMNS

Tips for Achieving Compliance

Last month, George covered practical EMC requirements for electronic instruments and systems. He wraps up the series with advice for achieving electromagnetic compliance.

By George Novacek (Canada)

I'll conclude this series of articles by touching on some the design aspects that an engineer must address to achieve electromagnetic compliance (EMC). As I explained in Part 1, a fluctuating voltage is always accompanied by a related current and vice versa, generating a fluctuating electromagnetic (EM) field. The magnetic component of the field and permeability (μ) associated with it cause magnetic flux H . Electric component and permittivity (ϵ) associated with it cause electrical flux. In consequence, every fluctuating voltage or current causes a propagating EM wave. Resistivity (R) in the path of the electrical current causes energy loss which converts into heat. There is always a return current into a circuit node identical in every way to the transmitted current out of that node, but in opposite polarity.

In the far field, EM power is divided among all alternative paths, inversely proportional to the impedance of each path. The wave impedance through vacuum or air, associated with radiated emissions and immunity, is approximately 377Ω . If an equipment is exposed to near field, determination of the impedance is much more complicated and characteristics of the field must be taken into account.

EMI PREVENTION

One objective of good EMC design is to prevent electromagnetic interference (EMI) emanating from our equipment, while making the equipment immune to external sources of EMI. Both requirements go hand in hand, often utilizing the same methods of suppression.

For effective EMC, it must be addressed early in the design stage. If left unattended for later time, what could have been a very simple solution, such as moving a few components on the printed circuit board (PCB), could become a nightmare, costing thousands in design modifications, retesting, and schedule slip.

The first step to EMC is determination of the requirements. Most electronic equipment is subject to compliance to some standard. Even if there are no mandatory requirements, responsible engineers should not be satisfied with passing just the minimum standard. Your engineering pride should not allow you to put your name to a design that could be inadvertently affected by a cell phone operating in the vicinity or a similar embarrassing flaw.

The first line of defense is a competent PCB layout and proper grounding. Keep the traces carrying high frequencies short, shield components known to be susceptible to or




PHOTO 1
A metal cabinet is an excellent shield

generating noise. Ground, not necessarily synonymous with “Earth,” is a sink for interference currents to be directed back to their source. PCB ground layers, separate for analog and digital circuits, are very effective. If only a single or two-layer PCB is used, ground bars are an effective solution with the added benefit of acting as stiffeners and perhaps a heat sink.

A metal enclosure is an effective EMI shield (see **Photo 1**). Such an enclosure, milled out of a solid piece of aluminum, is not always affordable. Sheet metal or fiberglass with an embedded mesh are commonly used for industrial equipment. Consumer equipment is often in plastic cabinets, sometimes with an

internal metallic coat for shielding.

EXTERNAL INTERFACES

Before we delve into the details of protection circuits, let’s consider external interfaces, from power to loads to input devices, such as sensors. These are not always under the equipment designer’s control, but knowing their characteristics is crucial for effective EMC. With the exception of power-carrying conductors, all signal lines should be shielded. Double-braided twisted pairs, if affordable, are the best. Twists create small adjacent loops with opposite magnetic flux, which is very effective for canceling out the effects of the M-field.

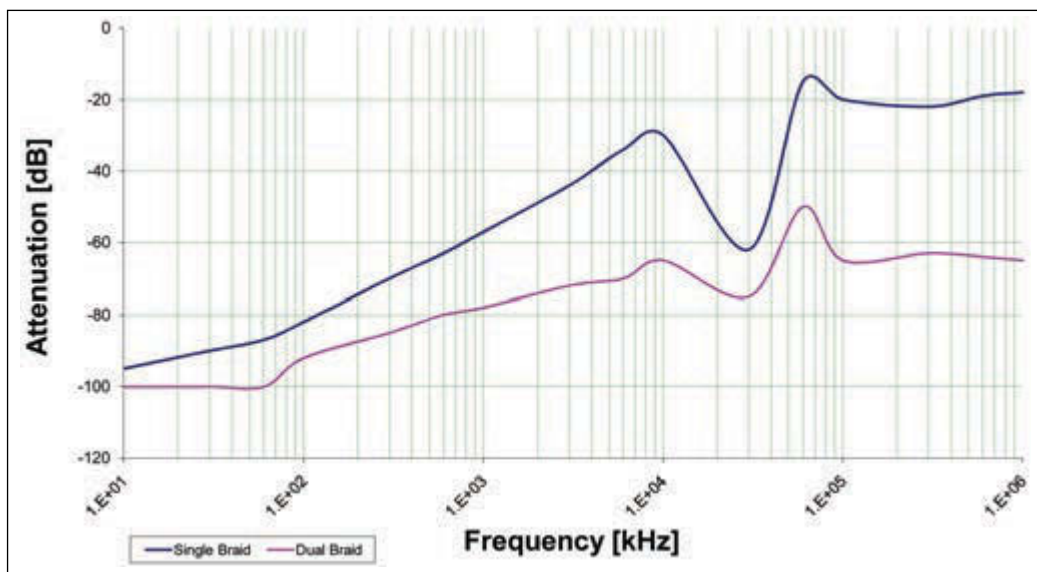


FIGURE 1
Shielding effectiveness of a single and a double braided cable




FIGURE 2

Typical EMI and transient suppression circuit

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.




PHOTO 2

Two stage EMI plate



circuitcellar.com/ccmaterials

G. Novacek, "Are You Grounded?," *Circuit Cellar* 145, 2002.

——, "My Analog World," *Circuit Cellar* 244, 2010.

——, "The Shocking Truth About EMC," *Circuit Cellar* 117, 2000.

RESOURCE

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

RTCA, "Environmental Conditions and Test Procedures for Airborne Equipment," DO-160G, 2010.

Nonferrous cabinets or shields, often made of aluminum or copper, effectively suppresses the E-field, but are ineffective for suppression of the M-field. That is best achieved with high permeability ferrous materials, such as mu-metal.

Figure 1 shows the shielding effectiveness of a single and a double braided cable irradiated by E-field in a laboratory. Notice that at high frequencies a single braided cable provides less than -20 -dB attenuation. Shield terminations and the use of connector shells have crucial effect on the shielding effectiveness. The popular pigtail shield termination, simple and quick to make, acts as an inductor, degrading the shielding effectiveness. The effectiveness can be quickly improved by employing the 360° shield termination instead. The connector shells, if not metal or metalized and connected to the shield, have a detrimental effect on shielding, as they cause discontinuity of the shielding.

There is a common dilemma whether to ground the shield at one end only or at more points. The general rule of thumb for lower frequencies, such as analog (not RF) signals, is to ground the end at the controller only. The shield, however, must be bonded to the shell of a sensor assembly, but may be floating with regard to the system chassis ground. For high frequencies, such as to suppress radiated emissions caused by digital circuits, bonding at both ends of a cable and even at several points in between, may be needed to satisfy the certifying requirements.

Knowing the maximum field strengths we have to work with, we can calculate the interference levels reaching inside the device, we can measure the interference generated by our circuits and determine the maximum levels the circuits can accept without performance degradation. **Figure 2** shows a typical transient suppression and EMI filtering to use in extreme environments. It is located in a grounded metallic cavity. This dual stage circuit can be modified for less demanding applications. The first stage is a spark gap. In **Photo 2** the spark gaps are within the yellow rectangle. Many are dual gaps, meaning that if one fires the second one fires too. They are most useful when protecting balanced circuits but may be used for separate circuits as well. The ones in the picture fire at 90 V and can conduct up to 5,000 A to ground. They have 52-V DC holdover, meaning that the DC supply must be less than that to allow the gaps to extinguish the discharge.

Sometimes the gap can be omitted, using just the transzorb (TVS diode) together with R1 to limit a spike. R1 and the transzorb are also seen in Photo 2, just next to the gaps. Typical value of R1 is 10 to 30 Ω . For DC power




PHOTO 3

Controller PCB with dirty and clean cavities

supplies and loads chokes are used instead to minimize the DC voltage drop. The TVS diodes act like fast, high dissipation Zener diodes. They are unidirectional as well as bidirectional for balanced or AC signals.

Due to the finite speed of gaps and TVSs, there is always a remnant, narrow, high-voltage spike, also caused by an electrostatic discharge (ESD). Capacitor C1 together with the following low pass filter suppress it. The feed-through filters leading into the clean cavity with the electronics are seen at the extreme right of the EMI plate.

The filters are usually "pi" filters with a -60 dB/decade theoretical slope. They're manufactured with critical frequency from about 10kHz to a few megahertz.

Connectors with filter and TVS pins are sometimes sufficient to achieve EMC, but they do have limitations. They can withstand relatively low voltage spikes only, dissipate a limited power and be quite costly.


DESIGNING FOR EXTREMES

In extreme environments, the enclosure is divided between electrically "dirty" cavity, such as shown in Photo 2, where cleaning up the signals takes place. Clean signals are conducted to the clean cavity through the filters.

The cabinet shown in Photo 1 is an example of a design for an extreme EM environment. The ridge in the bottom of the upper cavity marked by the black rectangle is bonded to the PCB ground plane which is marked by the yellow rectangle in **Photo 3**. Components to the left of that area are EMI suppression circuits. The controller circuits are to the right. This is a dual redundant controller with each channel occupying a vertical half of the cabinet.

Aside from the EMC method described above, great deal of immunity can be achieved with optically coupled interfaces. Common

mode signals can be effectively suppressed by transformers. And when you are dealing with E-fields of thousands of volts per meter, you can always employ fiber optics.

To conclude, don't forget that the EMC testing is static. In a real life, dynamic environment there is always a chance that a combination of several conditions can result in a failure. That's when the fun begins. 



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



ABOVE THE GROUND PLANE

Startup Transients

Ed provides insight into some startup problems arising in his retrofitted sewing machine controller, with solutions involving both hardware and software.

By Ed Nisley (US)



PHOTO 1

This image, taken through a visible-blocking infrared filter, emphasizes the IR LED on the proximity sensor board. Surprisingly, the black drive belt appears transparent in IR.

Perhaps you've proceeded all the way through a system's design, collected and assembled the parts, and written the bring-up code, only to discover the real challenge: making it work the way you expected at the beginning. The problems you encounter seem obvious only after they occur!

In this column, I'll discuss several parts of my sewing machine controller retrofit project that required more tweaking than I expected. Although your projects should never have any of these problems (because they're obvious!), perhaps my experiences can serve as a debugging guide when you're faced with a new and somewhat cantankerous collection of parts that should work perfectly together.

MOTOR SPEED & SHAFT POSITION SENSORS

The Arduino firmware must have feedback from the sewing machine to determine the motor speed and needle position, so I used a pair of optical proximity switches based on the TCRT5000 IR sensor and available for under \$2 on eBay. You can find modules with various sensor orientations, but they all include a sensitivity/range adjustment, a digital output bit, and an analog




PHOTO 2

The setscrew on the motor pulley produces a glitch as it passes over the phototransistor. Rather than attempting to smooth out the screw, I fixed the glitch in software.

voltage output. The infrared image of the motor RPM sensor in **Photo 1** highlights the IR LED on the right, with the phototransistor behind the barrier just to its left.

I blackened the motor pulley with Brownell's Oxpho-Blue Gun Bluing solution, which, despite the name, produces a dense black finish on steel that's a perfect "dark side." Stainless steel HVAC tape, burnished with a plastic rod, boosts reflectivity and increases the contrast of the other half. The sensitivity trimpot has enough range to turn the output off over the tape and on over the blued steel, so I set it

at the midpoint of those values for the best average noise immunity.

The interrupt handler should run once per motor revolution, triggered by the falling edge of the signal, when it records the elapsed time since the previous interrupt. The first version of my handler sometimes reported that the motor ran twice as fast as shown on the display of my trusty laser tachometer, which meant the interrupt happened twice on each revolution.

Unsurprisingly, at least in retrospect, the setscrew locking the pulley to the motor shaft

```

struct pulse_t {
    byte Counter;
    unsigned long LastIRQ;
    unsigned long Period;
    unsigned int RPM;
    byte State;
};

struct pulse_t Motor;
struct pulse_t Shaft;

//-----
// Motor RPM interrupt handler

void ISR_Motor(void) {

    static unsigned long Now;

    digitalWrite(PIN_SYNC,HIGH);

    Now = micros();

    // detect and ignore glitches
    if ((5000ul < (Now - Motor.LastIRQ)) && !digitalRead(PIN_MOTOR_REV) ) {
        Motor.Counter++;
        Motor.Period = Now - Motor.LastIRQ;
        Motor.LastIRQ = Now;
        Motor.State = digitalRead(PIN_MOTOR_REV);    // always zero in a Physics 1 world
    }

    digitalWrite(PIN_SYNC,LOW);
    return;
}

//-----
// In setup() function

pinMode(PIN_MOTOR_REV,INPUT_PULLUP);
pinMode(PIN_SHAFT_POSITION,INPUT_PULLUP);
attachInterrupt((PIN_MOTOR_REV - 2),ISR_Motor,FALLING);           // one IRQ / motor rev
attachInterrupt((PIN_SHAFT_POSITION - 2),ISR_Shaft,CHANGE);     // two IRQs / shaft rev

```

LISTING 1

Each falling edge from the motor RPM sensor triggers this interrupt handler, which should run once per motor revolution. However, the setscrew holding the pulley to the shaft can produce spurious interrupts, which the the if() statement detects and ignores. The machine shaft sensor interrupt handler uses similar logic, so it's not shown here, but the last two lines configure it to run on both edges of the sensor signal.

PHOTO 3

This position of the counterweight on the sewing machine's main shaft indicates that the needle is at the top of its travel. The needle will be at the bottom after half a revolution, with the counterweight's other edge over the sensor.



produces a glitch on the rising edge of the sensor's digital output, as shown in the upper

trace of **Photo 2**. The interrupt handler records the time between successive interrupts and that glitch introduced a second interrupt on each revolution. Just to add confusion, certain motor speeds produced two glitches, while other speeds produced none at all.

Triggering the interrupt handler on the rising edge would produce a second interrupt shortly after the first, which didn't seem helpful. I was reluctant to grind the screw flush with the pulley and put tape over it, as I must occasionally remove the pulley for adjustments. As a result, I decided to fix this one in software.


The interrupt handler in **Listing 1** marks its execution by turning on an output bit while it runs. That bit appears in the lower trace of Photo 2, where it shows the Arduino hardware and interrupt header code requires about 7 μ s to save the CPU registers before passing control to the handler; the first cursor should mark the falling edge of the glitch, not the rising edge.

The handler's `if()` statement accepts only those interrupts that occur more than 5 μ s after the previous one, with a low voltage at the sensor pin. That logic works well over the normal range of motor speeds, although I can still trigger a spurious interrupt by hand-turning the motor very slowly.

With that sensor working, I moved on to the shaft position sensor, shown in **Photo 3** with the shaft turned to set the sewing machine needle at top dead center. The shaft rotates counterclockwise and the counterweight occupies half of the circumference, producing

FIGURE 1

Plotting the setpoint current, open-loop actual current, and shaft speed against motor speed shows the variation that the firmware must process. The difference between the setpoint and actual current depends strongly on transistor temperature.



a square wave with rising edges aligned at the needle's highest position and falling edges at its lowest position.

Unlike the motor speed sensor, the position sensor produces a clean signal. The strip of stainless steel tape brightens the counterweight's machined surface enough that the sensor can easily distinguish it from the cutaway section.

The shaft position interrupt handler resembles the motor speed handler, minus the setscrew glitch rejection logic. The last two lines of Listing 1 show the `setup()` code that configures the handler to run on both the rising and falling edges of the signal, when it records the state of the sensor after each edge. That bit indicates the needle position at the time of the interrupt: low = needle down, high = needle up.

The two sensors and their interrupt handlers produce accurate speed measurements, but they don't directly indicate that the motor has stopped turning. A few lines in the main loop detect that the motor speed handler hasn't run in the previous 100 ms: the absence of that interrupt means the motor isn't turning.

Plotting the shaft speed against the motor speed produces the blue points in **Figure 1**. The vertical scattering below 500 RPM shows where the motor stalls at low speeds, without




PHOTO 4

The upper trace shows the motor current limited to 650 mA. The firmware samples the current just before the falling edge of the first nine pulses in the lower trace, then sorts the values during the last pulse.

the stall detection code. The plot includes several consecutive runs with increasing and decreasing speeds that produced a maximum spread of 6%, sufficiently accurate for a machine controlled by a foot pedal. An eyeball-fit line has a slope of 0.089, so the drive belts

```
#define NUM_I_SAMPLES          9

unsigned int SampleCurrent(byte PinNum) {

    unsigned int Samples[NUM_I_SAMPLES];
    unsigned int AvgSample;
    byte i,j;
    for (i=0; i < NUM_I_SAMPLES; i++) {        // collect samples
        Samples[i] = ReadAI(PinNum);
        delayMicroseconds(640);
    }
    for (i=0; i < (NUM_I_SAMPLES - 1); i++)
        for (j=0 ; j < (NUM_I_SAMPLES - 1 - i); j++)
            if (Samples[j] < Samples[j+1]) {
                Samples[j] ^= Samples[j+1];    // swap entries!
                Samples[j+1] ^= Samples[j];
                Samples[j] ^= Samples[j+1];
            }

    AvgSample = (Samples[1] + Samples[2])/2;    // discard highest sample

    return AvgSample;
}
```

LISTING 2

This function samples the values from the Hall effect sensor monitoring the motor current and returns an estimate of the maximum value. The highest reading generally corresponds to a transient, so the code averages the next two lower values.

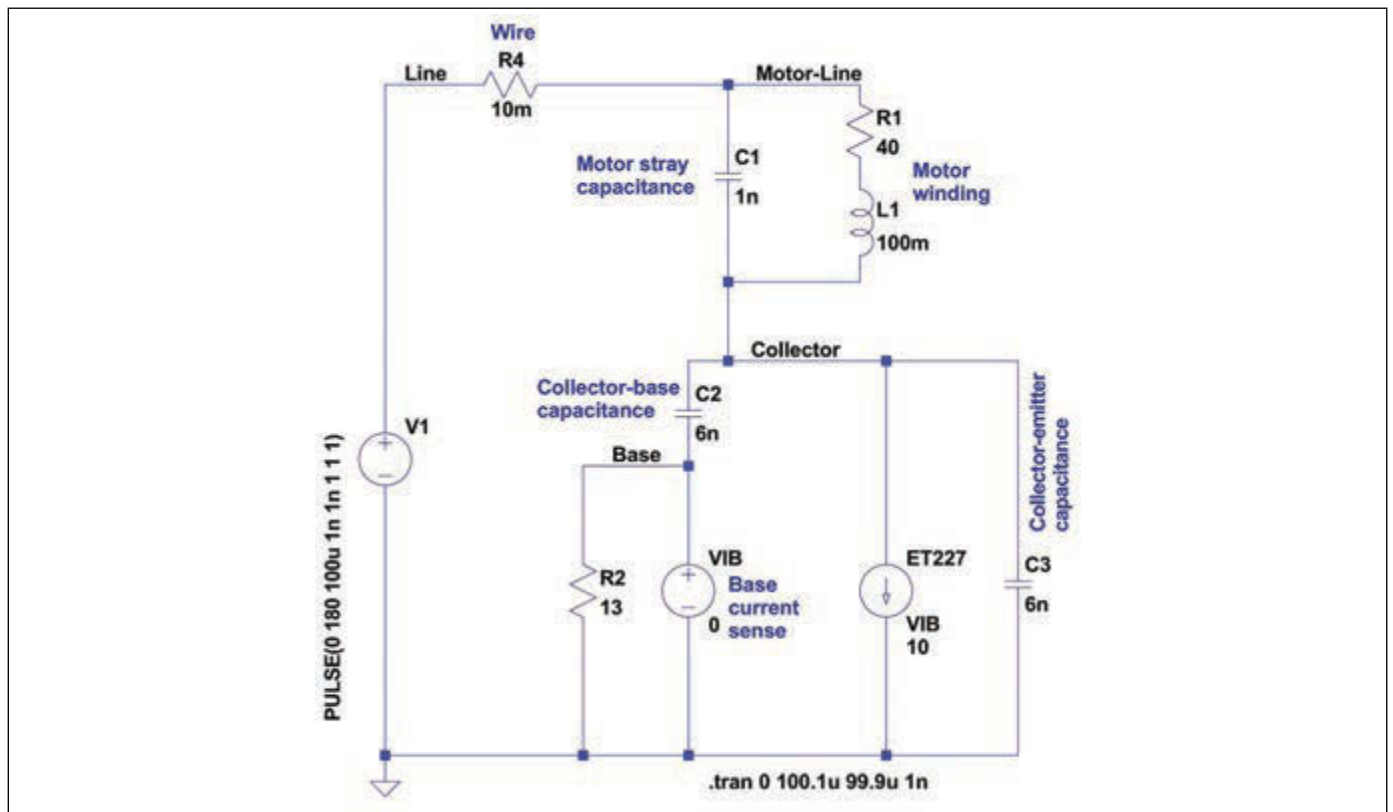


FIGURE 2

A simple SPICE model of the motor and ET227 transistor just after the relay contacts close. All values represent estimates based on measurements, but the overall accuracy won't be very good.

reduce the motor speed by $1/0.089 = 11.3$. The other points scattered around that plot require a bit more explanation, though.

CURRENT SAMPLING

As you saw in the January 2015 column, the motor controller uses a big ET227 NPN transistor as a current limiter, with a bridge rectifier converting the AC supply voltage into unfiltered DC. Even though the universal wound motor runs perfectly well from the half-sine voltage, the oddly shaped current waveform in Photo 4 shows the effect of the current limiter. As the current limit increases, the flat tops shrink as the waveform becomes a distorted sinusoid, with the motor inductance holding the current above zero between the pulses.

I had originally planned to synchronize the current sampling routine with a zero-crossing voltage detector, but the current waveform peak doesn't occur at a consistent delay from the voltage minimum. The lower trace in Photo 4 shows a different approach: collect nine measurements spanning one cycle of the rectified waveform (one half of a complete power-line cycle) into an array, then apply digital signal processing to determine the maximum current. Even through the

measurements aren't synchronized with the power line, their results include data representing the entire waveform.

Because the leading edge of the flat tops in the current-limited waveform generally have a small overshoot, as shown in Photo 4, the code in Listing 2 sorts the array in descending order and returns the average of the second and third measurements. That work well for most cases, but will slightly underestimate the maximum of a waveform without flat tops. Although it's not complex DSP, it gets the job done with enough accuracy for my purposes.

The calibration routine steps the motor current limit from 400 to 900 mA in 50 mA steps, records the ADC value returned by the code in Listing 2, then dumps those values to the Arduino's serial port. I also manually record the "true" motor current, as measured by my calibrated Tektronix Hall effect current probe, at each current step. Plotting the true current against the ADC count produced a line so straight there's no point in showing the graph:

$$\text{Current (mA)} = 2.97 \times \text{ADC} + 152.6$$

That equation applies for ADC values from

SUPERIOR **EMBEDDED** SOLUTIONS



DESIGN YOUR SOLUTION TODAY
CALL 480-837-5200

www.embeddedARM.com


NEW! **TS-4900 Computer Module**
1 GHz i.MX6 with WiFi & Bluetooth



Pricing Starts At
\$99
Qty 100
\$134
Qty 1

- 1 GHz Single or Quad Core Cortex A9 ARM CPU
- Up to 2 GB DDR3 RAM and 4 GB eMMC Flash
- Wireless 802.11 b/g/n and Bluetooth 4.0
- 4 KLut FPGA, 1x Gigabit Ethernet, 1x PCI Express Bus
- 1x mSD slot, 1x SATA II, 1x USB Host, 1x USB OTG
- 70x DIO, 4x I2C, 1x I2S, 2x SPI, 2x CAN
- 1x LVDS & 1x RGB Display Port
- -40 °C to 85 °C Industrial Temperature Range
- Runs Android, QNX, Linux 3.10, Yocto, QT, OpenGL
- **Coming Soon:** Windows Support

NEW! **TS-7670 Industrial Computer**
GPS Radio and Cellular Modem



Pricing starts at
\$129 **\$168**
Qty 100 Qty 1



low cost plastic enclosure available

Features:

- 454MHz ARM CPU
- Up to 256MB RAM
- 2GB NAND/eMMC Flash
- 2x mSD Card Socket
- 1x 10/100 Ethernet
- 1x USB Host
- 4x DIO, 2x CAN
- 2x COM, 1x RS-485
- 1x Battery Backed RTC
- 1x Temperature Sensor

Benefits:


- Low power with 10mW sleep state
- -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

Coming Soon:

TS-7680: Like the TS-7670 w/ WiFi & Bluetooth



COMING SOON! **TS-7970: SBC Version of the TS-4900**



Available w/ TS-8550 Development Baseboard



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

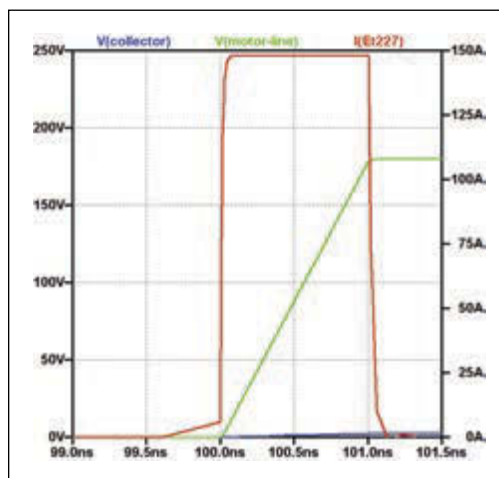


Unique embedded solutions add value for our customers

www.embeddedARM.com

FIGURE 3

The SPICE simulation shows that the transistor current can reach a high value for a very short time, driven by the parasitic capacitance between the collector and base junctions. The values depend strongly on the voltage rise time caused by the closing relay contacts, which will be entirely different from that shown in the model.



200 through 360, corresponding to motor currents from 400 through 900 mA. The gain coefficient accounts for everything from the number of turns on the ferrite core and its permeability, to the Hall effect sensor and differential amplifier gains. The offset value includes the differential amplifier's manual offset, which I set to about 100 mV to keep the output well above 0 V with 0 mA through the motor.

Note that the actual current limit set by the DAC at each step doesn't matter, because both the ADC and the Tek probe sample the same current. The equation simply converts from the ADC reading to the actual current, whatever it may be.

THE CLICK OF DOOM

Collecting all of that data required many test runs, repeatedly starting and stopping

the motor under control of the Arduino firmware. As one test started, the relay clicked closed, the motor accelerated to full speed, and nothing I could do affected the speed. Some careful probing revealed that the DAC still produced the expected series of voltages and the current amplifier drove the correct current through the optoisolator LED, but the big ET227 NPN transistor wasn't paying any attention.

More probing showed that the ET227 had failed with a short circuit between the collector and emitter terminals. In a typical application using the transistor as an on-off switch, a collector-emitter short can cause serious collateral damage until the fuse or circuit breaker pops. In this case, however, the motor simply ran with a direct connection to the AC power line and everything else survived.

I decided the ET227 probably failed from simple age, because it came from a decades-old stash that my buddy Eks had accumulated, and installed a spare. Two days later that "new" transistor failed with a collector-emitter short under exactly the same conditions. Obviously, I missed an important clue.

The Safe Operating Area graph in my January column showed that, even though the ET227 can control 100 A and 1 kV, the sewing machine motor operates just inside the continuous SOA region at 1 A and 180 V. Faced with two dead transistors on my desk, I built the SPICE model in **Figure 2** to see what could push the transistor's operating point over the edge into failure.

Keep in mind that a simulation cannot prove that a design *must* work. At best, it can suggest why a circuit *may* fail, by allowing you to model conditions that you cannot create and measure in physical hardware.

None of the transistors in the SPICE library came close to the ET227's Herculean ratings, so I used an idealized model: a current-controlled current source driven by the transistor's base current. All the internal resistances will be very small in a transistor rated for 100 A continuous current and I felt that omitting the forward biased base junction wouldn't affect the results.

Because the failure was intermittent, I thought that it might occur only when the relay closed at the peak of the AC power line waveform, applying an abrupt high voltage pulse to the motor and transistor. Although the inductive motor winding wouldn't transmit an instantaneous current change, the 1 nF of stray capacitance I measured between the windings and the case implies a similar capacitance between the winding terminals.



circuitcellar.com/ccmaterials

RESOURCES

E. Nisley, "Brute-force Motor Control," *Circuit Cellar* 294, 2015.

——, "ET227 Transistor: SOA Violation," 2014, <http://softsolder.com/2014/11/18/et227-transistor-soa-violation/>.

——, "Kenmore 158: Shaft Position Sensor," 2014, <http://softsolder.com/2014/11/03/>

kenmore-158-shaft-position-sensor/.

——, "NTC 2.5 Power Thermistor Characteristics," 2014, <http://softsolder.com/2014/11/24/ntc-2-5-power-thermistor-characteristics/>.

——, "Sewing Machine RPM Sensing: Gun Bluing FTW!," 2014, <http://softsolder.com/2014/10/21/sewing-machine-rpm-sensing-gun-bluing-ftw/>.

——, "Universal Motor Control vs. Transistor SOA," *Circuit Cellar* 292, 2014

SOURCES

Brownell's Oxpho-Blue Gun Bluing Brownells | www.brownells.com

TCRT5000 IR Sensor Vishay Intertechnology | www.vishay.com
Optical proximity switch: search eBay for "tcrt5000 module"

The ET227 datasheet does not specify the transistor's junction capacitance values, so the 6 nF I measured between the collector and the other terminals with zero DC bias came as a surprise. I cracked the dead transistors open and found the huge silicon slabs bonded to the collector had enough surface area to account for that much capacitance. Of course, you should measure the junction capacitance using a differential bridge under proper bias conditions to get a more accurate value, a process that requires instrumentation I don't have.

Even without a SPICE simulator, the schematic suggests that the peak collector voltage won't be much more than 1/12 of the input voltage: the three capacitors form a series-parallel voltage divider across the input. Assuming a 180 V peak on the input, the collector terminal might reach 15 V, which certainly isn't enough to cause any problems by itself.

However, the collector-base capacitance could supply a base current that, amplified by the transistor's gain, would produce a damaging collector current. I thought that the motor's stray capacitance could apply a sufficiently fast voltage transient at the collector to drive current through the collector-base capacitance, and that, even with a low collector voltage and the internal 13 Ω base shunt resistor, the transistor's gain would produce enough collector current to force the operating point beyond the 100 A DC and 200 A pulse limits at the of the safe operating area plot.

A simple differential equation relates capacitor current to the voltage change at its terminals:

$$I_c = C \frac{dV}{dt}$$

Although mercury-wetted relay contacts can have a rise time on the order of a few picoseconds, I couldn't find any data on the dry contacts used in power relays. For lack of better data, I slowed the rise time by a factor of 1000 and **Figure 3** shows the results of the source voltage rising from 0 V to 180 V in 1 ns: a nearly rectangular 150 A pulse of collector current. The collector voltage remains under 3 V during the entire pulse and the current does its damage at essentially zero bias.

Assuming a faster rise time of 100 ps would increase the current to 1.5 kA, while a slower 10-ns ramp drives only 15 A through the collector. The actual rise time and peak current will depend on the relay contact rise time, as well as the stray series inductance that's not included in the model. Remember that the motor winding inductance doesn't

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.




affect the rise time, because the stray capacitance causing this problem completely bypasses it. The ET227 datasheet doesn't mention the effect of extremely short pulses, but the absolute maximum pulse current limit is 200 A at 50 μs.

Tinkering with the model showed that a small series resistor dramatically reduced the peak current, so I inserted four 2.5 Ω NTC power thermistors between the relay and the bridge rectifier. These thermistors, also known as inrush limiters, generally appear in power supplies, reducing the initial charging current of large electrolytic capacitors. Their resistance has a negative temperature coefficient, so the resistance falls as they warm up, reaching a steady state that depends on the power supply load.

The sewing machine motor will run intermittently, allowing the thermistors to cool and increase their resistance. However, even after a few minutes at the motor's full-throttle current of 1 A, each thermistor adds 1 Ω to the circuit and the quartet limits the simulated peak current to 33 A, well within the safe operating area.

As it turned out, the third ET227 has survived several weeks of relay clicks applied through the thermistors, so I'll say that problem has been fixed. I can't defend the simulation's exact results, although it does seem to reveal a failure mode that I'd completely overlooked.

CONTACT RELEASE

After the November 2015 column appeared, reader Steve Troy sent me a lengthy note describing a super-slow-motion AC motor drive he designed for a radar antenna mount. His technique should overcome the terrible low-speed torque of the sewing machine motor, without requiring complex control circuitry. I'll adapt his AC drive to the pulsed-DC circuitry you've seen here and report back: a simple matter of software should make the sewing machine work even better than I planned. Which is why I appreciate hearing from you: I learn something new every time! 

FROM THE BENCH

Solid-State Lighting (Part 2)

Get Your Lighting Under Control

In the first article in this series, Jeff introduced the topic of electroluminescent (EL) technology and its uses. This month he shows how EL technology can produce a nondirectional output as opposed to an LED's concentrated point source output.

COLUMNS

By Jeff Bachiochi (US)

Electroluminescent (EL) technology doesn't have the clout to challenge LEDs in replacing our heat-generating incandescent lighting. EL just doesn't have the output levels that can compete. What it does have is the ability to produce a rather exceptional nondirectional output as opposed to the concentrated point source output of an LED. Other than the high-voltage source required to get those photons flowing, it is relatively inexpensive to produce and easy to work with.

In the first part of this article series, I spent much of the column explaining the background of this material and what's required to make use of it. I suggest you refer back to that article as a primer if you happened upon this EL discussion midstream. I left off with a circuit that enables me to control eight separate pieces of EL material from a single high-voltage source that's available from many of the same vendors that offer EL in sheets, strips, and wire. Because EL material uses a high-voltage source, control is handled via logic level-driven TRIACs. The circuit uses a microcontroller to provide the TRIAC control of each channel. It also uses a serial user interface that enables you to set up control sequences for each channel. Each channel's sequence consists of

a combination of functions with corresponding durations. Four functions are offered: Channel On, Channel Off, Brighten, and Dim. The circuit I presented last month used a subset of these, On and Off. Dim and Brighten were not provided.

HIGH VOLTAGE

By itself, the term "high voltage" generally denotes danger. This warning should be carried forward when experimenting with EL. Last month I provided a schematic that showed an isolation of the high-voltage circuitry from the logic-level circuitry. It is a good idea to keep these physically separated from one another so you can keep the two both electrically and physically isolated (see **Photo 1**). Putting this all into an enclosure will keep fingers and toes away from danger; however, while on the bench, it is especially important to recognize the potential for harm.

One of the issues with using a purchased high-voltage supply is you have no control over it. That is, it will free run at a frequency that may make it difficult to control. The circuit presented uses TRIAC drivers that are either phase-independent or phase-dependent. Phase-independent control enables the TRIAC



to turn on at any point in its cycle, while phase-dependant control limits when a TRIAC can turn on (only at zero crossings). Syncing control to zero crossings enables one to consistently turn on a TRIAC at the same point every cycle. Without this sync, you have On/Off control but no way of varying this over each cycle.

Providing “per-cycle” control over our 60-Hz world is easy because we know what the frequency of each cycle will be and we can monitor the voltage to get “in-sync” with it. We know each cycle will be 1/60 of a second, or 16.6 ms. To operate a TRIAC at the 50% level, the TRIAC ON signal must delayed by 4.1 ms after each zero crossing. Most of these high-voltage circuits run at a soft frequency and you can’t depend on them to remain constant. It becomes very difficult to monitor the frequency and dynamically adjust delays based on their varying zero-crossing times.

Readers who are familiar with *Star Trek*, and specifically “The Wrath of Khan,” may be familiar with Captain Kirk’s solution to the Kobayashi Maru. He defeated the so-called no-win situation by hacking the problem and providing a “back-door” if you will. To start off this month, I am going to eliminate a variable and thus allow a solution to present itself.

Based on inverter topologies, an addition to the present circuit includes a microcontrolled oscillator with complementary outputs driving a transformer’s primary (in this case, the step down transformer is reversed) to produce 120 V

in its secondary. A 12-V source is applied to the transformer’s center tap and alternately driven FETs ground the winding (see **Figure 1**). The FET’s PWM source will then control over the period of time during each half cycle that the FET is turned on. This moves the phase control from each TRIAC to this oscillator. The transformer’s secondary can now produce a voltage from 0 to 120 V. By syncing this with a channel’s on time, we can produce a brightening and dimming effect for each channel.

PWM

There are couple of ways to get PWM outputs that will alternate on separate pins. The first might be to use the top pair of a full-bridge output. This approach requires changing direction every cycle to alternate direction (outputs). A second approach might be to use the top and bottom of a half-bridge configuration. Again you must change direction every cycle to alternate between the outputs. I chose to use a third approach using a single PWM output. This approach makes use of an




PHOTO 1
I kept the high-voltage circuitry at one end of this prototype. There was plenty of room left over to add the circuitry for an on-board DC-to-AC converter to last month’s project.

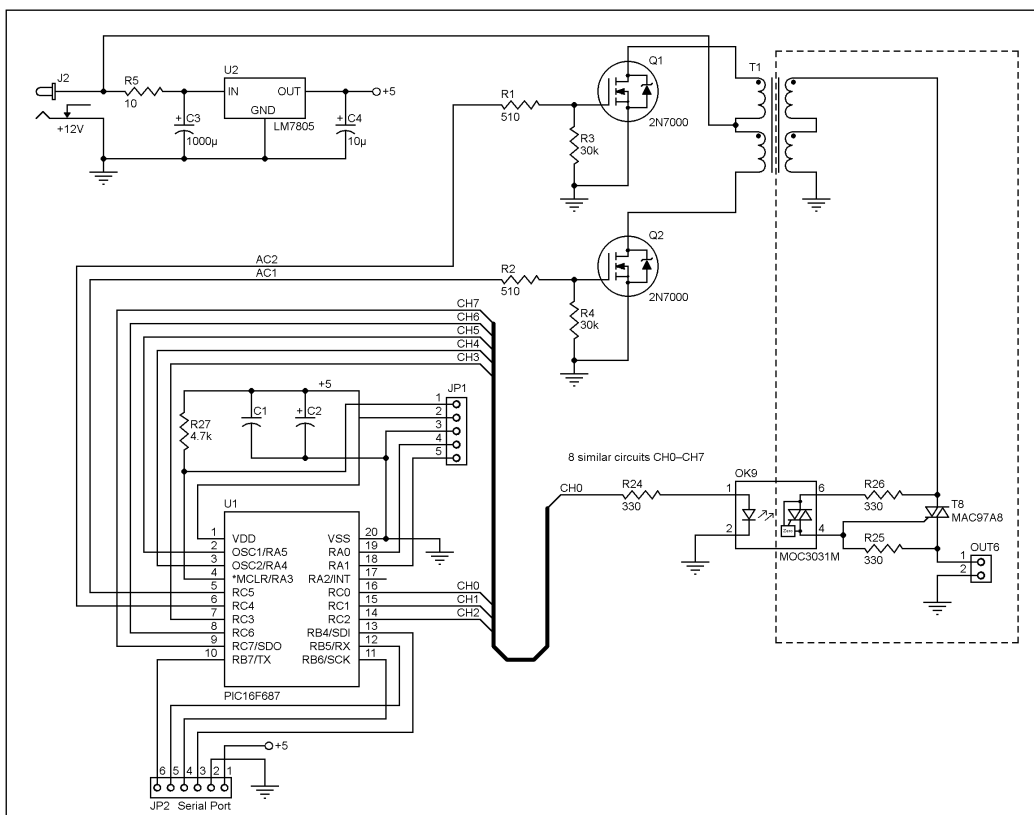


FIGURE 1
The right side of this schematic shows the circuitry I added to last month’s project. The microcontroller’s hardware PWM output is “steered” to drive each FET alternately. The frequency on the primary is therefore 2x the PWM rate. A 100% duty cycle is maximum RMS secondary voltage. A 0% duty cycle is zero RMS secondary voltage. So, the duty cycle controls the RMS secondary voltage.

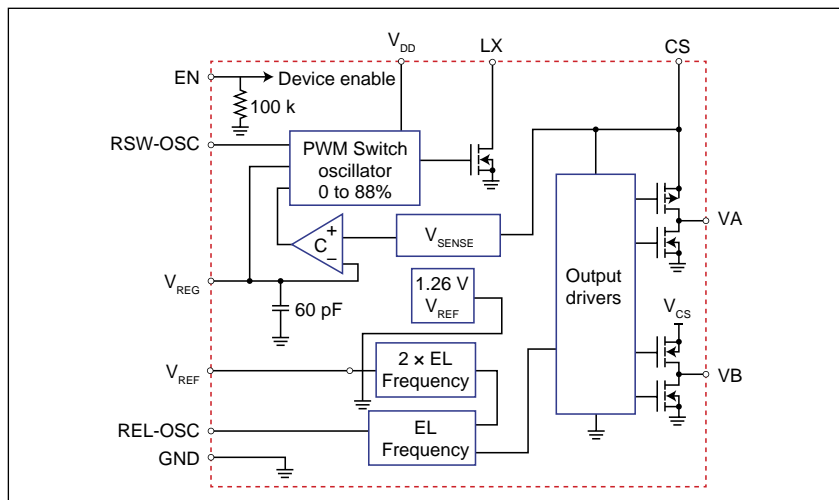


FIGURE 2

This block diagram shows the internal reference voltage that is fed into the V_{REG} input to control the inductor/diode tied to LX to pump up the diode/capacitor on CS to a regulated output voltage of 110 V. The output drivers then use the voltage on CS to produce ± 110 V to an EL lamp connected between VA and VB.

enhanced function called Steering mode. A single PWM output can be routed to any (or all) of four output pins via bits in the steering register. An interrupt available at the end of each PWM timer overflow is used to trigger an interrupt to swap output pins.

Two timers are used in this example. Timer 4 provides a tick every 100 ms for the Brighten and Dim commands. This is the command's minimum time period. Every 100 ms, all durations are decremented. Should a duration reach zero, the next command for that channel is queued up. Timer 6 sets the PWM frequency with the PR6 register. The PR6 register is used to reload Timer 6 each time it reaches zero, providing the base frequency for the PWM output. The CCP1L register will hold the PWM's on time pulse count. It is some fraction of the PR6 register. If you place a value of half PR6 in this register, the PWM output will be on for half the PWM period and then turn off for the second half.

While both outputs are never really on at the same time, the FETs these outputs drive could potentially be on at the same time if they turn on faster than they turn off. In a bridge circuit, if both FETs are on at the same time, it creates a high current path across the power supply. This is not exactly the case here, but it would cause unwanted harmonics here. To eliminate this possibility, you can add a delay to the on time in each cycle. The lower seven bits of the PWM1CON register indicates how many ticks the PWM ON should be delayed.

Since the outputs alternate, the actual output frequency will be half the PWM frequency. While the resistance of the transformer will limit the DC current through the transformer and FET, it is a good idea to completely remove any drive to either FET when the oscillator is stopped. Under program control, you can do this by disabling the drive outputs. However, you should use caution when running an in-circuit debugger, as you may inadvertently break the program while the oscillator is on, which would create a static condition.

Last month, I used a simple "Eat at Joe's" scenario, where the circuit sequenced through commands for the first three channels of EL control. The commands instructed the first channel turn on then off (Eat), followed by the second channel (At), and finally the third channel (Joe's) remaining on for 3 s.

```
C1 05 X5 X5 X5 X10 X10 X10 X10 X10 <CR>
C2 X5 X5 05 X5 X10 X10 X10 X10 X10 <CR>
C3 X5 X5 X5 X5 010 010 010 X10 X10 <CR>
```

Using this same idea we can substitute Brighten and Dim commands here to get a fade effect instead of blink of each EL segment.

```
C1 B5 D5 X5 X5 X10 X10 X10 X10 X10 <CR>
C2 X5 X5 B5 D5 X10 X10 X10 X10 X10 <CR>
C3 X5 X5 X5 X5 B10 010 D10 X10 X10 <CR>
```

DURATION TO PULSE WIDTH

Now that there is a way to alter the PWM that drives the step-up transformer and you can adjust the secondary's AC output voltage, you need to be able to vary the pulse width during each duration period. As I explained earlier, the CCP1L register controls the on time of each PWM period. We will need to plug values in here that range from 0 (0%) to PR6 (100%). I use the duration value for this.

The duration value decrements once each 100 ms to zero. Let's begin with the Dim command because it directly relates to this decreasing value. In the above example, the D5 command calls for a Dim to occur over five 100-ms periods. So, you use the following formula to calculate a value for the pulse-width register CCP1L: $PR6 \times \text{PresentDuration} / \text{CommandDuration}$. This calculation is used every 100 ms. For this command, you have the following five calculations.

$$\frac{208 \times 4}{5} = 166; \left(\frac{166}{208} = 80\% \right)$$

$$\frac{208 \times 3}{5} = 124; \left(\frac{124}{208} = 60\% \right)$$

$$\frac{208 \times 2}{5} = 83; \left(\frac{83}{208} = 40\% \right)$$

$$\frac{208 \times 1}{5} = 41; \left(\frac{41}{208} = 20\% \right)$$

$$\frac{208 \times 0}{5} = 0; \left(\frac{0}{208} = 0\% \right)$$

Every 100 ms, the duty cycle will be changed affecting the AC output voltage and the EL intensity. The Brighten command uses the same formula; however, instead of using the decreasing duration count, we subtract it from the Command Duration (in this case = 5) to get an increasing value and in turn a increasing AC output voltage.



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 or at www.imaginethatnow.com.

EL LAMP DRIVER

“And now for something completely different.” When Monty Python used this phrase, it usually meant more of the same British humor. While the project’s objective remains the same, I want to offer a different approach. So far, the DC to AC conversion has gone from a purchased module to an on-board converter. Both approaches are similar and use step-up transformers. I’d like to turn your attention to a family of parts manufactured by Supertex (now Microchip Technology). The HV8xx parts are EL lamp drivers. In particular, the HV860 part is a dimmable EL driver in a QFN-12 package. While the part uses an inductor, it is considerably smaller than using a transformer. And, since it has its own DC-AC converter, a connected EL device does not share the supply with other channels. In other words, all channels can be on at the same time.

Figure 2 depicts the HV860. Two internal oscillators use external resistors to set their frequency. The first oscillator drives an internal FET that allows the external inductor to charge and release its charge through a diode into an external capacitor on CS. The capacitor’s voltage on CS is monitored by V_{SENSE} and a portion compared to a 1.26-V voltage reference. The comparator’s output is feedback (duty cycle control) to the first oscillator regulating the voltage on CS.

The second oscillator converts the DC voltage on CS to an AC voltage on outputs VA and VB by driving four FETs in a full-bridge configuration. The converter’s EN input can be used to turn on and off the internal circuitry. However, since the internal reference voltage is connected to the voltage regulation comparator through an external resistor, you can modulate this via an external signal to reduce the voltage to the comparator. A lower reference voltage will in turn lower the DC voltage produced at CS. This effectively gives control over the AC voltage and thus the EL device’s illumination level.

There are a couple of limits for this device. First, it is designed to drive up to 5 square inches of EL material. Second, the external PWM signal used to adjust the reference voltage must be between 20 and 100 kHz. The PWM frequency would not be an issue for a microcontroller’s hardware PWM handling a single HV860; however, I am looking to handle eight devices with a single microcontroller. Since there are not eight hardware PWMs on many devices, this must be done in software. I used the same basic circuit from the previous example and just changed the high-voltage portion (see **Figure 3**).

TIMING

A 20-kHz output requires a cyclic output of 50 μ s. To adjust a pulse width of 0 to 100%




FIGURE 3

Here is the schematic of one channel of the eight EL drivers attached to PORTC of the PIC16F1829.

in increments of 1% would require resolutions of 500 ns. Even with this microcontroller’s PLL running at its maximum speed (Microchip Technology PIC16F1829) of 32 MHz, that’s an execution cycle of 125 ns (i.e., $1/32,000,000 \times 4$). There would only be four instruction cycles per bit of resolution required. This won’t fly. By limiting the PWM resolution to 10%, the number of instructions would skyrocket to 40. Even this seemed doubtful, but let’s give it a go.

First of all, we need to set up a timer to give 5- μ s interrupts. Refer to the flowchart in **Figure 4** to see what must be accomplished during Timer 2’s 5- μ s interrupt. Note that this

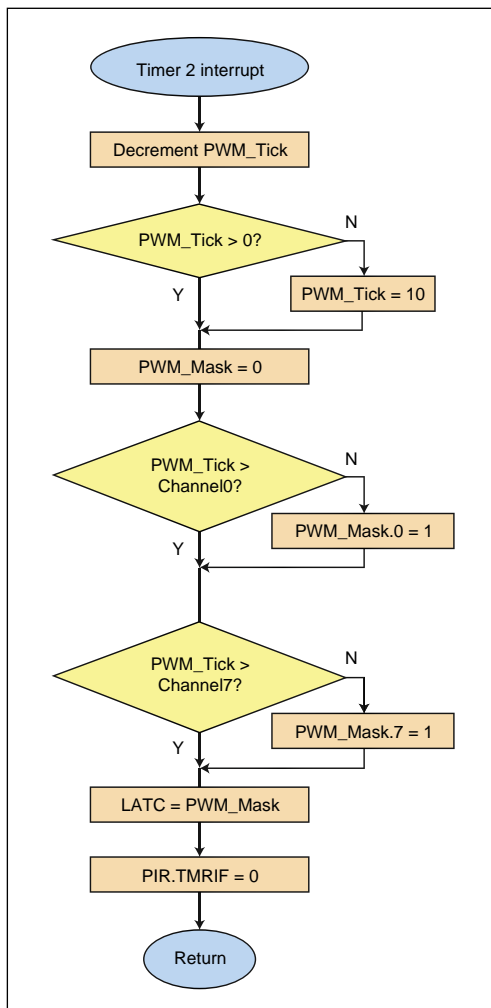


FIGURE 4

The flowchart shows what must happen every 5 μ s in order to produce eight channels of 20-kHz PWM outputs that have individual duty cycles adjustable to 10% resolutions. Note that channels 1 to 6 are identical and not shown here for simplicity.

COLUMNS

PHOTO 2

I find that working with SMT parts usually requires the high precision of a real PCB. Fine motor skills, tweezers, and a fine-tipped solder iron are a must. When ever possible, I try to use all through hole parts which is much easier for readers to duplicate.




diagram does not show channels 1 through 6, which are similar to channel 0 and 7 shown. I added some debug code to turn on an output bit at the entry to the interrupt and off at the exit so I could see its time as a positive pulse happening every interrupt. When I looked at this on the oscilloscope, I saw a 66% duty cycle and thought: "That's great, I've still got 33% of the execution time for other purposes."

What I didn't realize at the time was that the interrupt routine was taking just over 5 μ s, long enough for the interrupt timer to roll over 5 μ s again. The clearing of the interrupt at the end of the routine was in fact clearing the second interrupt that had just occurred and so it was another 5 μ s until it rolled over again. I could prove this by clearing the interrupt right away instead of waiting until the end of the interrupt. With that little change, everything stopped. This allowed the interrupt to reoccur before the end of routine, which kept interrupting and no other code could execute. This should be able to be fixed by lengthening the time between interrupts from 5 to 10 μ s. And sure enough, the interrupt timer was now happy to release control for the extra 5 μ s. However, I had just proven that I was not going to meet the 20-kHz modulation for eight output channels with the present code. I did find that even if I removed any debugging code, I could only lower the interrupt time to approximately 7 μ s before I began to see the

other functions begin to suffer. So, 8 μ s had to do. Now I'll cover what else had to happen between interrupts.

MAIN

The main loop code is similar to what had been already established. There are only three items to look for. First, the RXBuffer is checked for a received character. This only happens if the user is connected serially and enters commands to change the Animation Program. The default hardcoded Animation Program always executes on power-up unless the user enters a new one. That is then saved to internal EEPROM and becomes the new default. Second, the soft reset command flag is checked to see if the user wants to restart execution after entering a new Animation Program without powering down and back up again. Finally, a check of the execution flag indicates if at least one channel has decremented its duration to zero and is in need of a new command/duration pair.

For last month's example, I was concerned with enabling/disabling each channel. There was no control of the AC voltage because it was created by an external power supply module and I had no control of it. This month began with the addition of an on-board DC-to-AC converter. Besides enabling and disabling each channel, I created a PWM output. And thanks to a step-up transformer, I had control of the AC voltage by varying the PWM. If the change in PWM was synchronized with an enabled channel, I then controlled the AC voltage applied to that channel (for Brighten and Dim functions). In either case, only one channel was enabled at a time. Multiple enabled channels would put multiple loads on the AC and each channel would then share the current and illuminate at a reduced level (or not at all).

With this last example, each HV860 creates its own AC. That output is controlled via one of eight software PWM outputs generated for each channel. This requires an open collector type of output on PORTC, such that the output can sink to ground but will not source to VCC (3.3 V), as this would affect the voltage on the VREG input to the HV860. This microcontroller has push-pull outputs, so we need to be a bit tricky to simulate this. A logic 0 is written to the PORTC latch that handles sinking a channel to ground (off). When a channel is going to be on, the TRISC register is used to change the pin's function from an output to an input (a high-Z condition).

When an output is at ground potential (or off), none of the 1.26-V VREF can get through to the HV860's VREG input. When a PWM value of 0x00 (0% of PR2) is written to the duty cycle count, Channelx, the channel will always be off. When the port becomes an input (or on), it has



circuitcellar.com/ccmaterials

SOURCES

EL Panels, Strips, and Wire
www.adafruit.com | www.robotshop.com |
www.seeedstudio.com | www.sparkfun.com

PIC16F1829 Microcontroller and Supertex
 HV860 low-noise, dimmable EL lamp driver
 Microchip Technology | www.microchip.com


no effect on the 1.26 V VREF, and so it remains fully applied to the VREG input of the HV860. When a PWM value of 0xD0 (100% of PR2) is written to the duty cycle count, Channelx, the channel is always on. PWM values of anything in between 0x00–0xD0 will vary the average voltage on VREF. If the duty cycle is increased from 0% to 100%, the 1.26 V VREF will begin shorted out and then will be applied to VREG for longer periods as the duty cycle increases, until this modulation is at 100%, at which time VREG will equal VREF.

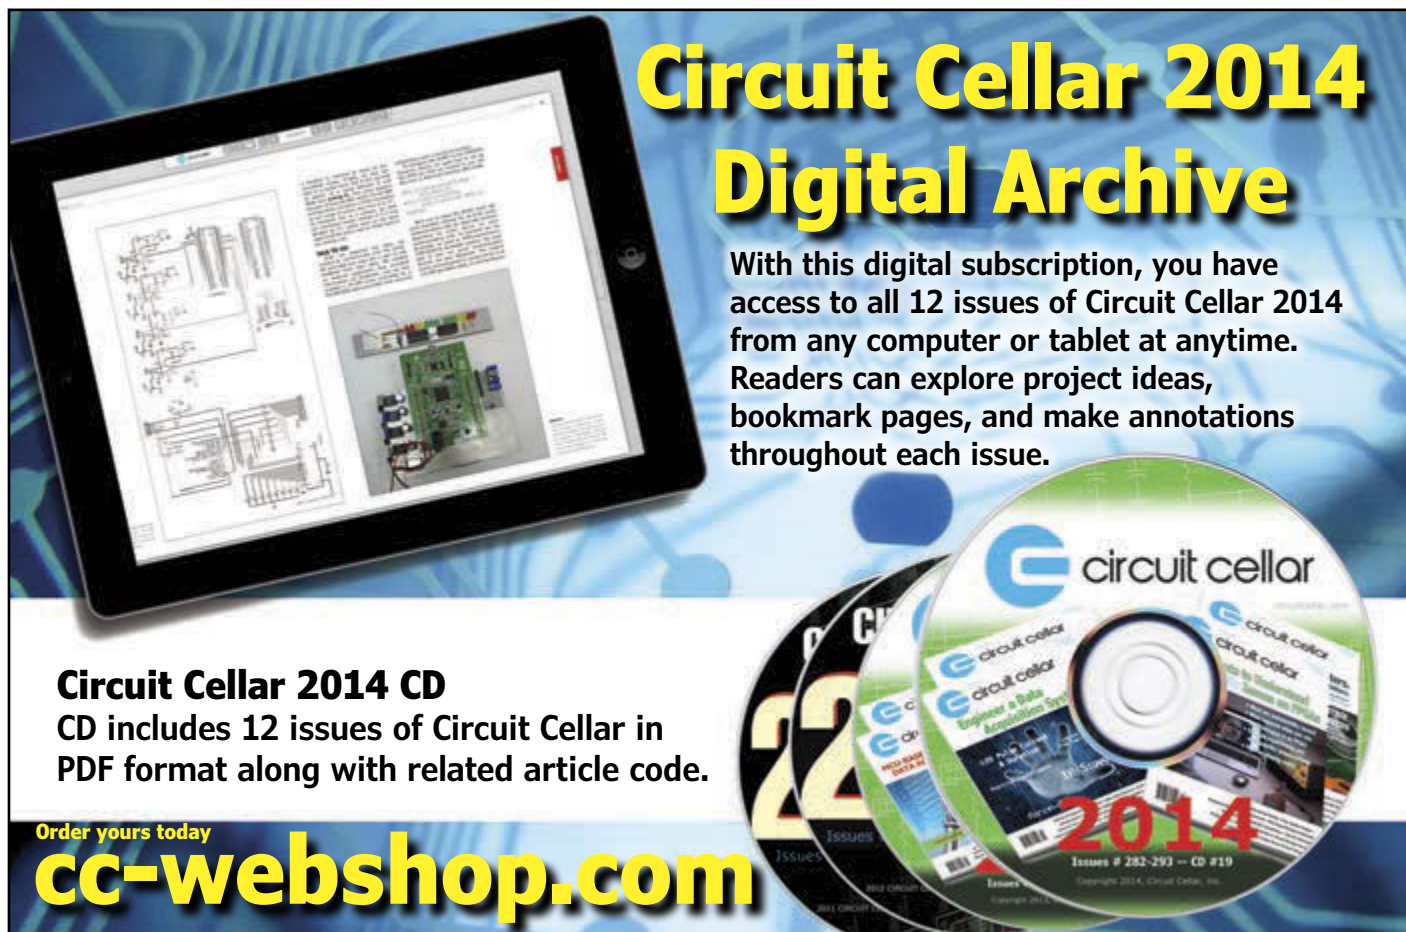
LIMITATIONS & ADVANTAGES

The HV860 is a tiny SMT QFN-12 part that's difficult to work with (see **Photo 2**). I used a toaster oven to mount these eight parts on the PCB before adding any other components. It's tough to get just the right amount of solder paste applied to each one, so they required a bit of touch up afterwards to remove a few bridges. My smallest iron tip, 0.4 mm, was just about right for the part's 0.5-mm lead spacing. But I'll have to wear magnification and use a steady hand if I want to successfully rework anything. The rest of the external SMT parts are huge by comparison once you've done the ICs.

The HV860 has a square pad, centered between the leads, that serves as a heat

transfer pad. The internal FETs can raise the operating temperature a bit. While the PWM signals I was creating are slightly below the minimum specs, the circuit seemed to operate well. However, if I wanted to be within specs, I was not going to get there with this microcontroller. There are other alternatives in the PIC18 family that offer a few advantages over the PIC16 device. These can execute at a faster speed (48 MHz vs. 32 MHz) and also offer two level interrupts. When you are counting cycles, every last one is important and you waste an amazing amount of support cycles just branching through the interrupt structure. When you have multi-level interrupts, you can assign an important task to the top level and therefore get right to it. Ten cycles might be the difference between attaining your goal, or not.

Each HV860 is able to drive 5 square inches of EL material. Check out the other parts available in this family. With each channel responsible for its own AC source, the animation program is free of restrictions. Now each channel can independently flash, brighten, and dim independent of the others. I'll bet a single channel programmable circuit could be pretty small if a tiny SMT microcontroller was used. I think I see the light! 



Circuit Cellar 2014 Digital Archive

With this digital subscription, you have access to all 12 issues of Circuit Cellar 2014 from any computer or tablet at anytime. Readers can explore project ideas, bookmark pages, and make annotations throughout each issue.

Circuit Cellar 2014 CD
CD includes 12 issues of Circuit Cellar in PDF format along with related article code.

Order yours today
cc-webshop.com

FROM THE ARCHIVES

Mini Rover 7

Electronic Compassing for Mobile Robotics

Electronic compassing is one of the most intelligent ways to provide absolute heading information for a mobile robot. In this article, Joseph explains why the PNI V2Xe compass turned out to be the best fit for his Mini Rover 7 robot, which he modeled after the NASA/JPL Rocky 7 Mars rover.

By Joseph Millier (US)

Recent technological growth has yielded some impressive tools for tackling increasingly difficult tasks. The robotics field has been a large beneficiary of these advancements because it encompasses so many disciplines, each heavily dependent on technology. Mobile robot navigation is one particular recipient of recent technological advancements. Accurate and reliable navigation is fundamental to the success of any mobile robotic application.

Navigation is typically given one of the highest, if not the top, consideration when designing an autonomous mobile robot. You have to consider numerous hardware and software options. These days a GPS receiver is one of the first tools that you think about to fulfill your navigational needs. However, the area that many mobile robots operate in is within GPS's present 3- to 15-m accuracy range.^[1] Even long-range mobile robots still have to maneuver successfully within that 3- to 15-m resolution uncertainty.

Compound these resolution limitations with satellite signal interference, and it becomes clear that a supplemental system is essential. Heading information, along with positional data from other sources, can provide interim position data as well as augment GPS data to resolve higher accuracies.

A magnetic compass has many advantages as a provider of heading information. Compassing is one of the only methods that can provide absolute heading information without external references for calibration. Today's electronic compasses easily interface with microcontrollers and come with a host of other features like low-power consumption and

built-in local distortion correction, such as any other instrument. Compasses have their own uncertainties and issues. Understanding how compasses work, as well as the behavior of the environment that they measure, will better prepare you to manage their uncertainties and best apply the technology to a project like my Mini Rover 7 robot.

BASIC ROBOT NAVIGATION

Dead reckoning (DR) is the fundamental navigation method used in mobile robotics. It is a method of mathematically tracking your present position by measuring speed and direction traveled at regular intervals, or distance and direction at any convenient interval. The latter is easier to do on most wheeled mobile robots if you use wheel encoders. This is also referred to as odometry.

The basic tools for DR in mobile robotics are a compass and an odometer. In more complex robots with advanced navigation resources, DR is still used to navigate between absolute position fixes. Position tracking by DR is used to get to and from fixed locations, help create maps of surroundings, and keep track of position when moving around unexpected obstacles.

Position errors have a tendency to accumulate over time when navigating with DR. A robot happily hobbling down a path can easily drift off the path because of the imperfect traction of its drive wheels or limbs. Furthermore, the terrain may be uneven, inconsistent, obstacle-ridden, or even move, as is the case for aquatic vehicles. A robot could simply correct its steering to stay on course (or on bearing), but that little off-course excursion could add a small position



error. The magnitude of the position error is the integral of time, distance, and off-course angle traveled while the robot is off course. What is really required is a new bearing calculation to increase the robot's chances to reach the desired target.

HEADING DETERMINATION

For centuries, we have been taking advantage of the Earth's magnetic field to orient ourselves with respect to its magnetic poles. Both the mechanical needle compass and electronic compass can provide absolute heading information. It is hard to beat the compass for this purpose. With the exception of GPS, other systems that exist require an external heading reference as calibration.

Gyroscopes use mechanical angular momentum changes to measure angular and linear movements. Traditional flywheel gyroscopes are fast-spinning gimballed flywheels with encoders. The encoders are situated about the pivotal axes of the gyroscope's gimbals and register angular movement of the spatially stable flywheels to its base, which is fastened to a host vessel. Modern gyroscopes use micro electromechanical systems (MEMS) and optical technologies in place of the bulky flywheels.

Gyroscopes have fast response times and are insensitive to magnetic anomalies. They are also relative angular position sensors, which require an external reference heading to initially set. A special kind of gyroscope called the gyrocompass can align itself with the Earth's rotational axis, but it tends to be a large and costly instrument.

Differential wheel encoding is another technique used to determine heading. Relative heading changes can be computed by taking the difference of distance traveled by two opposing wheels. This technique has the same traction and terrain issues associated with the aforementioned wheel encoder odometry.

A single-antenna GPS can provide heading information, but it is not instantaneous. It inherently lags the movement of the robot or vehicle because the derived heading requires previous position data. A GPS could not tell you where you are heading if you were to stop and change directions. Like compasses, GPS receivers do not require external reference heading calibration. Once moving, the GPS heading update rate is a maximum of approximately 1 Hz, although some receivers add damping, which increases this time constant even more. A dual-antenna GPS receiver can provide instantaneous heading—or yaw—information, although the recommended distance between the two antennas is 1 m. This fact, along with its large price tag, can be a limiting factor for many mobile robot applications.

A combination of techniques is the best




FIGURE 1
As you study the Earth's geodynamo magnet field pattern, note how the field lines are not horizontal to the surface except along the equator. The rotating dynamo and coil in the center represent the Earth's magnetic field being generated by its ever-flowing iron outer core.

approach. There are many ways to determine heading, each of which has its own strengths and weaknesses. None of them are infallible. For this reason, some systems use two or more methods cooperatively to increase system accuracy and reliability. The deciding factors are cost, accuracy, efficiency, features, availability, ease of use, speed, and size.

Kalman filters are typically used to integrate the data from multiple sensors to produce a more reliable and accurate heading. Kalman filtering is a statistical method that combines the dynamic model of the system with the statistical behavior of system errors. It enables navigation systems to handle periodic GPS signal interruption, odometer slippage, magnetic anomalies, and other sensor irregularities with minimal degradation of accuracy. Kalman filters also can be extremely complicated. You must fully understand the dynamic behavior of your systems and the statistical and systemic errors of your sensors in order to make proper use of Kalman filters. It might be easier and more feasible in less demanding projects to use other software-based analytical tools like averaging, weighted averaging, limiting, and majority voting to improve heading data reliability.

MAGNETIC COMPASSING

The Earth's magnetic field is created deep in its iron core by a regenerative magnetic field




FIGURE 2
I plotted the magnetometer sensor output versus the angle. You can also see the x-y plot of the magnetometer sensor output.

Measured sensor value		Quadrant(s)	Heading Calculation
X	Y		
≥ 0	> 0	270–360	$360 - \text{Arctan}(Y/X)$
≥ 0	≤ 0	0–90	$0 - \text{Arctan}(Y/X)$
< 0	All	90–270	$180 - \text{Arctan}(Y/X)$

TABLE 1

Use this table to convert the limited range, Cartesian coordinates system output of the Arctan function into full-range compass coordinates.

ABOUT THE AUTHOR

Joseph Miller is a Senior Electrical Engineer at PNI Corporation, where he designs consumer and OEM sensor- and network-based products. He earned an ASEE degree from the Heald Institute of Technology in San Francisco, and has been an electrical engineer for approximately 20 years. Joseph enjoys competing and demonstrating his robots. You may contact him at compass@sonomatron.com.

generator that's sometimes referred to as the geodynamo. This iron core has a liquid outer section and a solid inner section. The flow of electrical current in the turbulent liquid iron outer section creates the magnetic field.

The simplest description of the Earth's magnetic field spatial pattern is that of a dipolar field with magnetic flux emanating from the South Pole and converging at the North Pole. The Earth's magnetic field pattern is a little more complex than a simple bar magnet model. As previously mentioned, the Earth's geodynamo is constantly moving. Presently, the magnetic poles are tilted about 11° away from the geographical poles, and they are not at exactly at opposite sides of the world either. The magnetic North Pole is located in northeastern Canada, and the magnetic South Pole is located in the Antarctic Ocean south of Australia.

The geographical North Pole is also known as true north. The angular difference between the true poles and the magnetic poles at a given location is called the declination angle. Depending on your location, true north could appear to either the east or west of the magnetic North Pole.

The Earth's spherically shaped geodynamo produces a magnetic field as shown in **Figure 1**. Note that the field lines are not horizontal to the Earth's surface, except at the Earth's magnetic equator. Unlike the Earth's straight geographical equator, this one meanders but is located in roughly the same area. At the magnetic poles, the field lines are vertical. The angular vector of the magnetic field with respect to the horizontal plane at any given location is known as the dip angle, or inclination angle. The density of the magnetic field also varies around the world. The magnetic field density is approximately two times as dense at the magnetic poles as it is at the equator.

The magnetic field vector is sometimes referred to as having two separate components, a horizontal component and a vertical

component. At the magnetic equator where the magnetic field is horizontal, the field has no vertical component. At the magnetic poles, the field is purely vertical and has no horizontal component. At places where the inclination angle is 45° , the horizontal and vertical components are equal.

The U.S. Geological Service (USGS) and the Nation Oceanic and Atmosphere Administration (NOAA) maintain web sites that have global maps and on-line programs that chart declination angles, field intensities, dip angles, and much more. I will focus on the horizontal component of the magnetic field, because that is the portion that contains the heading information that I wish to measure.

MAGNETIC MEASUREMENTS

A magnetometer is an instrument that can measure the flux density of a magnetic field. It uses one of any number of types of sensors to convert magnetic flux to voltage, current, frequency, or some other electronically measurable form.

There are numerous types of magnetic field sensors: the saturable core magnetometer (or fluxgate magnetometer), the Hall effect sensor, the magneto-resistive sensor, and the magneto-inductive sensor. A two-axis magnetometer, in which the two sensors are in quadrature (orthogonal) orientation, can be used as an electronic compass to compute heading. When it is parallel with the measured field, the magnetometer sensor's output is at maximum for the given amount of magnetic flux density that is present. When the magnetometer sensor is perpendicular to the magnetic lines of flux, the sensor will output no signal. A plot of the x-sensor output versus the y-sensor output results in the heading being represented around the polar axis of the coordinate system origin (see **Figure 2**).

This form is preferred as a visual analysis tool for sensor and system performance analysis and troubleshooting. Notice that the y-axis is inverted from that of a typical Cartesian coordinate system. This was done so that the compass coordinates would be produced in its correct orientation. When operating with compass coordinates, it is important to remember to make the proper translations from a Cartesian coordinate system to a compass coordinate system, especially after using trigonometric functions.


At angles between parallel and antiparallel with respect to the magnetic lines of flux, the sensor's output signal, X, is a product of the applied magnetic flux density, b, and the cosine of the angle, q, of the sensor from being parallel with the flux lines.

$$X = \beta \cos(\theta)$$

[1]

FIGURE 3

The x-y plots show various magnetic distortion effects.



If a second sensor is added, and if it is positioned at a right angle to the first sensor, its output, Y, will have the same function as X, but will be 90° out of phase. The y sensor will be in the east position, and the x sensor will be in the north position. The two sensors are said to be in quadrature with one another. The equation for output Y is the following:

$$Y = \beta \sin(\theta) \tag{2}$$

You now have enough data to compute heading from the output values of the x and y sensors. Use the trigonometric identity:

$$\tan(\theta) = \frac{\sin(\theta)}{\cos(\theta)} \tag{3}$$

Combining Equations 1–3 yields:

$$\theta = \text{Arctan}\left(\frac{\beta \sin(\theta)}{\beta \cos(\theta)}\right)$$

or

$$\theta = \text{Arctan}\left(\frac{Y}{X}\right) \tag{4}$$

The arctangent, or inverse tangent (tan⁻¹), is inherently restricted to ±90°, which covers only two quadrants of the coordinate system over its entire input range of -∞ to ∞. This function also operates in the Cartesian coordinate system, which is rotated 90° from compass coordinates.

To convert the Cartesian q to a compass heading coordinate, a translation must be performed. **Table 1** shows the translation based on the output polarity of the sensors when the sensors are oriented as shown in Figure 2. Note that magnetometer polarities may vary depending on the manufacturer.

As demonstrated in Equations 1 and 2, the output signals of both sensors are proportional to field density (b) and its angle relative to magnetic north. The field density can be extracted at any angle of the quadrature pair by computing the geometric sum of the two sensors outputs:

$$\beta = \sqrt{X^2 + Y^2} \tag{5}$$

Equation 5 computes the horizontal component of the overall magnetic field density from the perspective of the robot’s horizontal plane. By monitoring this value, you can spot magnetic anomalies and tilting, which effect heading accuracies.

SOFT/HARD IRON DISTORTION

As you now, producing a compass heading with a pair of sensors is fairly straightforward. However, there’s trouble in the neighborhood. Magnetic anomalies are all around you. Other magnetic sources and alternate magnetic flux pathways are corrupting your intended signal—the Earth’s magnetic field.

Magnetic fields are generated by electrical




FIGURE 4
Here you can see the effects of soft iron distortions in the Earth’s magnetic field.

current (electromagnetism) in your circuitry and neighboring appliances. There are permanent magnets in your robot’s motors, and there is magnetized metal in the robot’s construction materials.

Items external to the robot, such as appliances, underground pipes, and ducts may have their own fields, and can locally offset or change the direction of the Earth’s field. Items that create offsetting fields are classified as hard iron. A soft iron material is anything that alters the natural path of a magnetic field’s lines of flux. By “natural path” I mean its path through air. Put another way, any material with a relative permeability other than one is a soft iron material. A soft iron also can have a hard iron effect. Usually soft iron and hard iron materials are referred to by the effect they have on the surrounding magnetic field. These effects are easy to recognize when an x-y graph is plotted of the compass sensors output values as the compass is rotated through 360°. **Figure 3** illustrates these effects.

All of the plots assume that the culprit material is local, which is to say that it is




FIGURE 5
Here are several magnetic signature plots of the Mini Rover 7 with additional effects.



PHOTO 1

No, it isn't the Spirit rover you've seen on the news; it's my Mini Rover 7 robot.

onboard the robot when it is rotated. These traces, or profiles, can be thought of as the magnetic signatures of the robot's local environment. Notice how the soft iron effects create an elliptical profile. Soft iron material attracts, bends, and condenses magnetic flux toward itself, creating sparse field densities on its lateral sides. At the same time, it will also create dense magnetic fields in areas that are in line with it and the source field direction, as shown in **Figure 4**. The distribution and quantity of soft iron material will determine the alignment and offset of the ellipse's foci.

Nearby electromagnetic fields, ferrous construction materials, and motors make robots especially tough environments for compasses. Before using an electronic compass in a mobile robot, it is a good idea to analyze its magnetic signature by creating an x-y sensor plot. From this plot, you can identify the existence of the aforementioned simple distortions and, more importantly, the more complex ones. The plots should give you an indication of the compass's chances of providing accurate readings, as well as information to help you make corrective adjustments to the robot or sensor placement.



PHOTO 2

I thought you'd like a look inside the Mini Rover 7's enclosure.

Modern electronic compasses have little trouble mathematically correcting for simple soft iron and hard iron distortions, sensor gain mismatch, and even sensor misalignment (orthogonality), which are all shown in **Figure 3**. Electronic compasses calculate the correctional coefficients after completing a one-time calibration procedure that involves rotating the robot or vehicle through at least one complete rotation so that it can analyze the geometry of the x-y sensor data.

When testing your robot and creating plots for your analysis, the patterns to watch for are excessive hard iron offsets and multimode soft iron distortions, which look like lumpy circles or ellipses. Excessive hard iron offsets can bring sensors into their nonlinear regions. This looks like a flat spot on the x-y circle plots, which are perpendicular to the axis of the saturated sensor. It is also a good idea to create plots for the robot when it is rotating both clockwise and counterclockwise, and then compare the two to see if the motor currents affect sensor readings.

Figure 5 shows x-y sensor plots (magnetic signature plots) running on my Mini Rover 7 robot. The red trace shows that the robot has some hard iron offsets but no soft iron distortions. I rotated the robot in both directions under its own power and saw nothing significantly different, although it would be hard to see any rotational offsets with these plots. There are ways to test this.

The slightly elliptical blue trace is the result of placing a 9-V battery next to the compass. I placed a small magnet in the robot when I generated the light blue trace. The dark blue trace is the result of placing a PC chassis next to the robot when it performed a rotation. The small amplitude is to be expected because the external soft iron material attracts the magnetic flux away from the robot, therefore creating an area with a sparse field, much like the one in **Figure 4**. You should not calibrate the compass in your robot under this condition because it does not represent a normal condition. Any electronic compass with hard iron and soft iron compensation should have no problem producing accurate heading data using any of the vehicle magnetic signatures shown in **Figure 5**.

As you have seen, with a little effort a compass can adapt itself to your robot's unique internal magnetic signature by using its built-in calibration algorithm. However, when operating after calibration, your robot will most likely come across external hard iron and soft iron objects, which will alter the Earth's magnetic field direction.

Figure 4 shows how a soft iron object can alter the heading of a robot as it attempts to keep steady (north in this case) when it moves past the object. Notice that the field density is sparse in this area. As long as the sensors are relatively close (so that they experience the same field density ratios), the compass abilities are not significantly diminished. This is because the heading calculation for a quadrature sensor pair uses the arctangent of the ratio of the two sensors (see **Equation 2**). The overall field density, or magnitude of the field density, can be monitored by calculating the geometric sum of the two sensors (see **Equation 4**), which can alert you to such anomalies. This information can be used to weigh your options.

You may decide to temporarily trust differential wheel encoder outputs for tracking heading while the overall measured field density falls outside established thresholds. A Kalman filter or your own statistical algorithm can use this information with other sensor data to make improved estimates of actual heading. It is important to use the compass's calibrated x-y sensor data to compute field density because it already has been compensated

SUPERIOR **EMBEDDED** SOLUTIONS



DESIGN YOUR SOLUTION TODAY
CALL 480-837-5200

www.embeddedARM.com

NEW! TS-7250-V2 Embedded Board High Performance and Industrial Grade



Pricing Starts At
\$165
Qty 100
\$199
Qty 1



Available with TS-ENC720 enclosure
(Shown with optional microSD card)


Features:

- Up to 1 GHz ARM CPU
- 512 MB RAM
- 8 or 17 KLut FPGA
- 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 half a second
- Easy development w/ Debian and Linux 3.14
- High data reliability with 2 GB SLC eMMC flash
- -40 °C to 85 °C industrial temperature range

TS-7700 Embedded Board High Performance Compact SBC



Pricing Starts At
\$119
Qty 100
\$165
Qty 1



Available with TS-ENC750 enclosure
and TS-752 interface board

- 800 MHz or 1 GHz Marvell PXA168 ARM CPU
- 512 MB DDR3 RAM, 128 bytes NVRAM
- 2x microSD Card Sockets for Redundant Storage
- 8 KLut FPGA, 1x 10/100 Ethernet, 2x USB Host
- 55x 3.3V DIO, 6x TTL UARTs, 1x I2C, 1x 12S, 1x CAN
- On board temperature sensor and watchdog timer
- -40 °C to 85 °C Industrial Temperature Range
- Runs Linux 2.6.34 or 3.14 with Debian Wheezy



Computer-on-Modules State of the Art Embedded Design

- TS-4600: 450 MHz low cost w/ 2 Ethernet
- TS-4710: Up to 1 GHz PXA168 w/ video
- TS-4720: Like TS-4710 + 2 GB eMMC Flash & 2 Ethernet
- TS-4200: Atmel ARM9 w/ super low power
- TS-4800: 800 MHz FreeScale iMX515 w/ video



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

www.embeddedARM.com

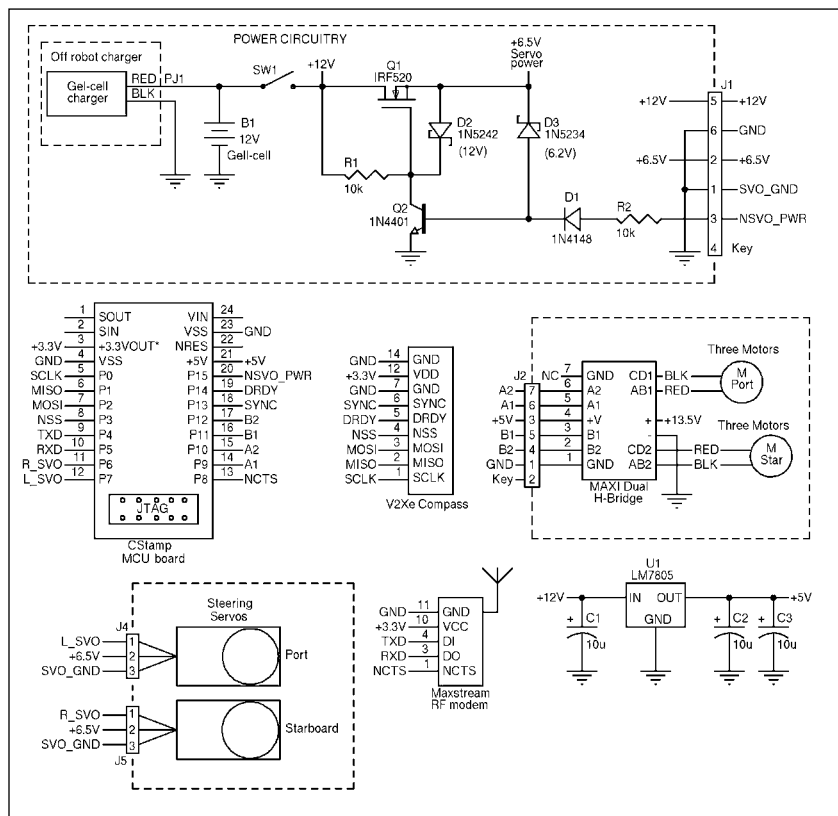


FIGURE 6

Because I like to make my robots as expandable and serviceable as possible, most subsystems are modularized and plug into the main board using connectors.



circuitcellar.com/ccmaterials

REFERENCES

[1] Garmin Ltd., www.garmin.com/aboutGPS/.

RESOURCES

H.R. Everett, *Sensors for Mobile Robotics: Theory and Applications*, A.K. Peters, Wellesley, MA, 1995.

R.B. Langley, "The Magnetic Compass and GPS," *GPS World*, September 1, 2003.

E. Kaplan ed., *Understanding GPS: Principles and Applications*, Artech House, Norwood, MA, 1996.

M. Miller et al., *The Personal Robot Navigator*, A.K. Peters,

Wellesley, MA, 1998.

PNI Corp., "Multipoint Calibration Primer," 1001766 R01, 2003.

U.S. Geological Service (USGS), geomag.usgs.gov.

SOURCES

CStamp Module
Ken Gott | cstamp51@earthlink.net

Maxstream Wireless modems
Maxstream | www.maxstream.net

MAXI Dual H-Bridge two-motor driver kit
Mondo-tronics, Inc. | www.robotstore.com

V2Xe Compass module
PNI Corp. | www.pnicorp.com

C8051F330 Microcontroller
Silicon Laboratories | www.silabs.com

for your robot's personal magnetic signature. Modern electronic compasses like the PNI V2Xe can report the overall field density in proportion to the field that it experienced at every heading angle during a calibration cycle so that local field distortions do not interfere with the assessment of external magnetic fields.

EFFECTS OF TILT

Tilting an electronic compass can create heading errors. When tilted, the sensors no longer receive the magnetic field in the proportions measured during calibration. Equations 1 and 2 characterize the sensor output value with respect to the magnetic field's angle to the sensor's orientation in the horizontal plane. Tilting the sensor would expose the sensor to portions of the Earth's vertical field component and reduce its exposure to the horizontal field component, which could either be an increase or decrease in field density depending on local inclination angles and direction of tilt.

Equation 6 gives the relationship of heading error versus tilt angle when a compass is tilted in the north-south rotational axis (also known as pitch):

$$\theta_{\text{ERR}} = \text{Arctan}(\sin\alpha \tan\phi) \quad [6]$$

where θ_{ERR} is the heading error caused by tilt, α is the pitch angle compass, and ϕ is the inclination angle of the Earth's magnetic field.

To use an example, if your compass were located in San Francisco, which has a magnetic field inclination angle of 61° , you could expect a heading error of up to 1.8° for every degree of pitch for the first 10° of tilt. Tilt in the east-west direction (roll) or any compound angle of the two tilt axes creates similar errors, although the worst-case errors could still be characterized by Equation 6 by simply substituting the pitch angle, α , with a tilt angle in any direction. Another accuracy-degrading factor would be the lack of hard iron and soft iron distortion correction in the pitch and roll axes. The two most common ways to make a compass insensitive to tilt is to mount a two-axis compass on a gimbal and to use a three-axis, tilt-compensated compass.

COMPASS FEATURES

A good compass should have high dynamic range sensors. The need to compensate large hard iron offsets is not uncommon, and can account for 75% to 90% of the sensor's operating range. A good compass will have an operating range that is at least four times the Earth's magnetic field density.

A compass also should be able to resolve the Earth's horizontal magnetic field component to 1:115 for 1° resolution and 1:1146 for 0.1° . These numbers don't account for soft iron distortions that effect sensor gains and hard iron offsets, which would increase these requirements.

A good compass should be temperature-compensated. Magnetic sensors have temperature dependencies like most sensors. Fortunately, the sensor's common temperature effects drop out of the heading calculations because the arctangent function's input value is a ratio of the sensor pair (y/x). Sensor nonlinear temperature effects, external hard iron offsets, and soft iron distortions do not share this temperature cancellation characteristic. Of course, a compass that has hard iron and soft iron compensation is essential. The necessity of tilt compensation depends on the robot's requirements and your budget.

V2XE COMPASS MODULE

I used the V2Xe compass in the Mini Rover 7 robot. This is a 1" square electronic compass module that uses an SPI interface as a means of communication. It consumes less than 3 mW of power, and has an output resolution of 0.01° with a heading accuracy of 2° . Its field measurement range is about 20 times that of the Earth's field, which means that it can operate with extremely large hard iron offsets that are common in robotic applications.

The V2Xe compass can be calibrated using one of two methods. Distortion-compensation coefficients along with declination settings are stored in nonvolatile memory. The V2Xe can provide raw sensor data and compensated field magnitude. It has an adjustable digital low-pass filter for heading.

A continuous calibration is the simplest calibration method to perform on the V2Xe. Send the calibration start command to the V2Xe to begin the calibration process, rotate the robot in one or two complete circles, and then send a calibration stop command to end the calibration process. After completing the calibration, you can retrieve the heading data from the V2Xe as necessary.

V2XE IN THE MINI ROVER 7

The Mini Rover 7 depicted in **Photo 1** is a working scale model of the NASA/JPL Rocky 7 Mars rover. It incorporates six drive motors, and it has a rocker-bogie suspension as well as a zero turn radius steering system that's capable of rotating the rover in place. The robot contains a Maxstream 900-MHz wireless modem for receiving commands and transmitting data to and from a terminal program on a host PC, which has a matching wireless modem.

The Mini Rover 7 also sports a C8051F330 microcontroller that's mounted on a CStamp, which is a Basic Stamp II socket-compatible module. It contains two regulators and a JTAG programming/debugger port connector. I used the free integrated development environment and evaluation C compiler supplied by Silicon Laboratories along with its EC2 serial-to-JTAG adapter. The wiring diagram for the robot is shown in **Figure 6**.

Photo 2 shows the placement of the compass in the Mini Rover 7. What is not clear from this photo is that I raised the compass to distance it from the motors and current-carrying wires in an effort to reduce distortion. Additionally, the wires are twisted in complimentary pairs to alternate the magnetic field direction and to minimize the distance of the radiated field generated by current flow; this effectively reduces offsets to the measured field.

The end result is illustrated by the red trace in Figure 5, which shows that the robot's magnetic signature has only a small hard iron offset. The compass can easily compensate for this simple offset mathematically after a calibration is performed.

To explore the compass's performance in various environments, I wrote a simple program to perform basic moves

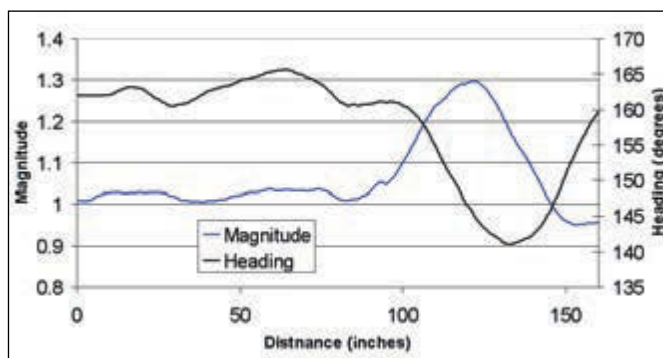


FIGURE 7

Here you have the robot's distance from the origin versus heading and magnitude during a straight-line test with a soft iron distortion effect near the 125" mark.

and transmit the robot's heading, as well as the measured field magnitude and raw sensor values. You may download the code from the Circuit Cellar ftp site.

The wireless modem that links the Mini Rover 7 to a remote terminal enables you to use the keyboard on the remote terminal to send the Mini Rover 7 commands for maneuvering and to select the type of data you want the robot to report. For instance, sending a "c" to the robot tells it to begin calibrating of the V2Xe compass. An "s" stops the calibration process and any movements. Sending an "h" tells it to include only heading data in its periodic data transmissions.


You can command the robot to rotate in a circle ("r" or "l"). It also can send raw sensor data ("b") in its datastream, which I was able to cut from a terminal program receive buffer window and paste into a spreadsheet program to create the x-y magnetic signature plots shown in Figure 5.

Another good test is to command the robot to move forward in a straight line. You can then observe the heading and magnitude values for irregularities. For this test, I commanded the Mini Rover 7 to travel down a hall in my house and into the living room. Along the way, the robot passed the kitchen. As it passed the backside of the refrigerator, the heading reported by the compass drifted. The magnitude was simultaneously changing, therefore signifying a magnetic field anomaly.

Figure 7 is a graph of the distance the robot traveled versus the heading and magnitude. Notice that the field magnitude is a good indicator of magnetic anomalies, which can be used to determine the quality of the compass heading information.

I performed an outdoor test that proved more successful than the indoor test. The outdoor environment didn't reveal any of the magnetic disturbances that I had observed indoors because I didn't encounter manmade iron objects.

THE NEXT MOVE

The program listing posted on the Circuit Cellar ftp site also contains a PID-based steering control system that keeps the robot on a steady course—if it doesn't encounter magnetic anomalies. The next step for this project is to add two wheel encoders to create a differential odometer. I plan to experiment with this data along with compass heading and magnitude information in a weighted averaging algorithm to increase heading accuracies. 

Editor's Note: This article first appeared in Circuit Cellar 165, April 2004.

CC SHOP



1

1 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



2

2 CC 2014 CD

2014 was an exciting year for electronics engineers! The continued success of open-source solutions, Internet of Things (IoT) revolutions, and green-energy consciousness has changed the face of embedded design indefinitely. In *Circuit Cellar's* 2014 archive CD, you can find all of these hot topics and gain insight into how experts, as well as your peers, are putting the newest technologies to the test. You'll have access to all articles, schematics, and source code published from January to December 2014.

Item #: CD-018-CC2014




3

3 MICROPROCESSOR DESIGN USING VERILOG HDL

After years of experience, Monte Dalrymple has compiled his knowledge of designing embedded architecture and microprocessors into one comprehensive guide for electronics engineers. *Microprocessor Design Using Verilog HDL* provides you with microarchitecture, writing in Verilog, Verilog HDL review, and coding style that enables you to depict, simulate, and synthesize an electronic design on your own.

Author: Monte Dalrymple

Item #: CC-BK-9780963013354



4

4 ADVANCED CONTROL ROBOTICS—NEW!

When it comes to robotics, the future is now! This book simplifies the theory and best practices of advanced robot technologies by providing insight to communication technologies, control robotics, embedded technology, programming, and more. Complete with code samples, schematics, and design tips. Great reading for beginners and experts alike.

Author: Hanno Sander

Item#BK-ELNL-978-0-963013-33-0

Further information and ordering: www.cc-webshop.com


CONTACT US: Circuit Cellar, Inc. | Phone: 860.289.0800 | E-mail: custservice@circuitcellar.com

What's your EQ? The answers are posted at www.circuitcellar.com/category/test-your-eq/. You can contact the quizmasters at eq@circuitcellar.com.

TEST YOUR EQ

Contributed by David Tweed

PROBLEM 1—Back in the days of stand-alone UART chips that required separate baud rate generators, one way to generate the 3.6864 MHz clock for the baud rate generator was to take the 16.000 MHz system clock and feed it to a synchronous 4-bit counter that was configured to divide by 13 by forcing it to load the value 3 when it got to 15, giving the following waveforms:



As you can see, the QB output of the counter produces 3 pulses for every 13 input clocks, and it turns out that this comes very close to the required frequency. What is the exact error, expressed as a percentage?

PROBLEM 2—Obviously, there is some jitter in the timing of the individual pulses produced by this circuit, relative to an evenly-spaced clock at the same frequency. What is the peak-to-peak magnitude of this jitter?

PROBLEM 3—Modern UARTs usually include internal baud rate generators that can divide the input clock by an arbitrary integer N. Given an input clock of 16.000 MHz, and assuming that the output of the baud rate generator needs to be 16x the actual baud rate, what is the highest standard baud rate for which the frequency error is no greater than that generated by the scheme above?

PROBLEM 4—What is the maximum frequency error that a baud rate generator for a UART can produce? Assumptions:

- 8N1 data format
- Error equally distributed between transmitter and receiver
- Generator output is 16x the baud rate

Sign up for the FREE Circuit Cellar Newsletter!

You'll receive electrical engineering tips, interesting electronics projects, embedded systems industry news, and exclusive product deals via e-mail to your inbox on a regular basis. If you're looking for essential electrical engineering-related information, we've got you covered: microcontroller-based projects, embedded development, programmable logic, wireless communications, robotics, analog techniques, embedded programming, and more!

Subscribe now to stay up-to-date with our latest content, news, and offers!


circuitcellar.com



CROSSWORD

MARCH 2015

The answers will be available at circuitcellar.com/crossword.



ACROSS

2. Field of 60- or 120-Hz magnetic or electrostatic energy
5. Scan of a range of frequencies
6. WW
13. Pot
14. Prefix that divides a unit by a billion
15. Exhibiting the same physical properties in all directions
17. One millionth of a meter
18. To change in value by a discrete step
19. RFI, EMI
20. Test instrument that can make several different measurements

DOWN

1. QC (two words)
3. A field of radiated energy
4. Magnetizable substance based on iron
7. Bending of energy waves as they move around or through an obstacle
8. Any connection between two electrical conductors
9. Slicing off of signal peaks
10. Conductor shared by various circuits
11. Slower than the speed of sound
12. Convert a digital signal back into an analog signal
16. Measured in fractions of a farad

IDEA BOX the directory of PRODUCTS & SERVICES

For current rates, deadlines, and more information contact Peter Wostrel at 978.281.7708 or circuitcellar@smmarketing.us.

Programming
 Debugging



\$89

For PIC® MCUs!

- ▶ ICD & ICSP functionality with CCS IDE Compilers
- ▶ Powers & Connects via USB
- ▶ FREE CCS Load Software
- ▶ Data Streaming with Real-time Diagnostic Data & more...

View our full line of Dev Tools:
www.ccsinfo.com/CC315
sales@ccsinfo.com
262-522-6500 x 35



PIC® MCU is a registered trademark of Microchip Technology Inc.

Adapters & Modules Turn-Key Solutions

Cost Effective High & Low Volume Solutions

- Device Replacement & Upgrades
- Subsystems & Daughter Cards
- Support for all IC packaging types
- Turn-Key Production & Assembly



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

MaxBotix®

High Performance Ultrasonic Rangefinders



Call Carlson today for application specific sensors

Phone: 218-454-0766
Email: info@maxbotix.com
www.maxbotix.com

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/ci.plc



TRIANGLE RESEARCH INTERNATIONAL

\$20 for 5PCBs

2 layer,4x4inch, FR4(RoHS),0.063",1oz, 2LPI,Green,1SK,Lead free HASL

Standard PCB: Promotion code: CC14061

PCB & PCBA

Small to Mass QTY

INSTANT QUOTE AT:
www.myopcb.com
OR CALL:
1-888-PCB-MYRO


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
Free 96 page catalog 1-800-826-5432



microEngineering Labs, Inc.

www.melabs.com 888-316-1753

Programmers for Microchip PIC® Microcontrollers



Starting at \$79.95

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

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.

The Future of Flexible Circuitry

By Mark Finstad



Mark Finstad is a Senior Application Engineer at Flexible Circuit Technologies in Minneapolis, MN. He is a nationally recognized expert in the design, fabrication, and test of flexible and rigid flex printed circuits with more than 30 years of experience in the flexible PCB industry.

The flexible circuit market has been growing steadily for the last three decades. This trend will continue into the foreseeable future as flexible circuitry supports many of the same industries and many of the same applications that have been around for more than 30 years. Past and current industries include military and avionics with most of these applications being high layer count, high-density rigid flex, and also consumer electronics, telecom, and automotive applications with flex circuit designs that are typically less complex than those of mil/aero. Medical diagnostic applications will continue to grow as new equipment is developed and older equipment is refurbished or redesigned. But if I had to sum up an answer to the question “where is flex going in the near future?” my answer would be simply “on you.”

The wearable electronics market has absolutely exploded in the last few years with new applications emerging almost daily. If an electronic device is going to be worn on the body comfortably, it has to be flexible. So what better way to provide interconnects for these types of devices than a flex circuit? Here are just a few of the current and emerging wearable products that contain flexible circuitry.

Wrist-Worn Activity and Body Function Monitors: Electronic watches were some of the first wearable

can be worn on the baby’s leg or in their clothing and can keep track of breathing, heartbeat, body temperature, etc. If the device senses that there is a problem, an alert is sent to the parents phone to wake them. It is almost like having a private nurse watching the child all night long.

Medical Sensors: This is an area that has been growing rapidly, and I predict that the trend will continue at an accelerated rate. With today’s push to get patients out of the hospital as quickly as possible, electronic home monitoring of the patient is going to be necessary. There are currently sensors that can be worn by the patient for several days at a time, while keeping tabs on heart functions continuously during this time. Just like the baby monitor referred to earlier, these devices can send alerts to the patient’s physician if any abnormalities are detected. These devices will allow a patient to recover from heart attack or surgery in the comfort of their own home while still having continuous monitoring of their state of health.

Pet Monitors: Even Rover gets to wear electronics these days. Training collars have been around for a while, but now thanks to shrinking electronics there are collars that contain GPS and mobile phone capabilities. Today a lost pet can use the GPS to figure out where he is and call his owner for a

If an electronic device is going to be worn on the body comfortably, it has to be flexible. So what better way to provide interconnects for these types of devices than a flex circuit?


electronics, so it was just a natural progression to include more advanced functionality than just time keeping. Wrist-worn activity monitors are light weight and use multiple axis accelerometers and other sensors to detect motion and body functions. They can capture and record daily activity levels as well as sleep cycles. This data is stored in on-board memory in the device until it can be downloaded to the user’s mobile phone. Since the human hand is larger than the wrist, these monitoring bands need to be able to expand when the user is putting it on or taking it off. Flexible circuits allow the band to flex while maintaining connectivity across flexing sections.

Foot-Worn Sensors: I have seen a lot of applications recently for electronics that are worn on the foot or inside the shoe. Foot-worn electronics monitor everything from steps taken when running or walking to stride irregularities that can contribute to back problems. These sensors need to be very thin in order to be comfortable and also very robust to survive in what I would consider a pretty hostile environment. Flexible circuitry is thin enough to lay on the sole of a shoe and be almost undetectable to the wearer.

Wearable Baby Monitors: Baby monitors are one of the newer products in the wearable electronics market. New parents no longer have to rely on a simple walkie talkie system to keep tabs on their little ones while they sleep. These monitors

ride home! Not really, but if your pet is wearing one of these devices he is never truly lost. The mobile phone module is used to transmit the GPS coordinates to tracking service, where the owner can log on and track the pet’s location to within a few feet.

Clothing Worn Electronics: This is an area that is just starting to emerge, and new technology is being developed to support these applications. Standard flex circuitry is constructed from a combination of polyimide film, thermo-setting film adhesive, and copper foil. Unfortunately, flex circuits fabricated with these materials will not survive the crumpling that they would be exposed to in a washing machine. I have seen several applications where flex has been incorporated into clothing that does not need to be machine washed (e.g., flexible heaters in winter gloves). The key to making this type of wearable application machine washable is to make the flex circuit not only flexible, but also stretchable. This means that both conductors and dielectrics must be developed that will allow the finished product to stretch and still maintain electrical continuity. This technology is not mainstream yet, but it is on its way.

These examples are just a small sampling of the applications that are currently on the market, and there are many others in development. As more and more of these applications emerge, flexible circuitry will continue to be the interconnect method of choice. 

CC Vault

Unlock the power of embedded design.



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.

A vault of need-to-know information in the fields of embedded hardware, embedded software, and computer applications

*CC Vault is a 16-GB USB drive.



Order yours today! cc-webshop.com

We bring the full range of *Electronic Contract Manufacturing services to your fingertip!*

FABRICATION



ASSEMBLY



KEYPADS



ENCLOSURES



This is the only place where you would put all your eggs in one basket to get fastest time to market. From concept design to prototype to full turnkey production on all your electronic products.



imagineering inc. www.PCBnet.com

**847-806-0003 sales@PCBnet.com
Certified Woman-Owned Small Business**