EDN europe

MARCH 2015

www.edn-europe.com

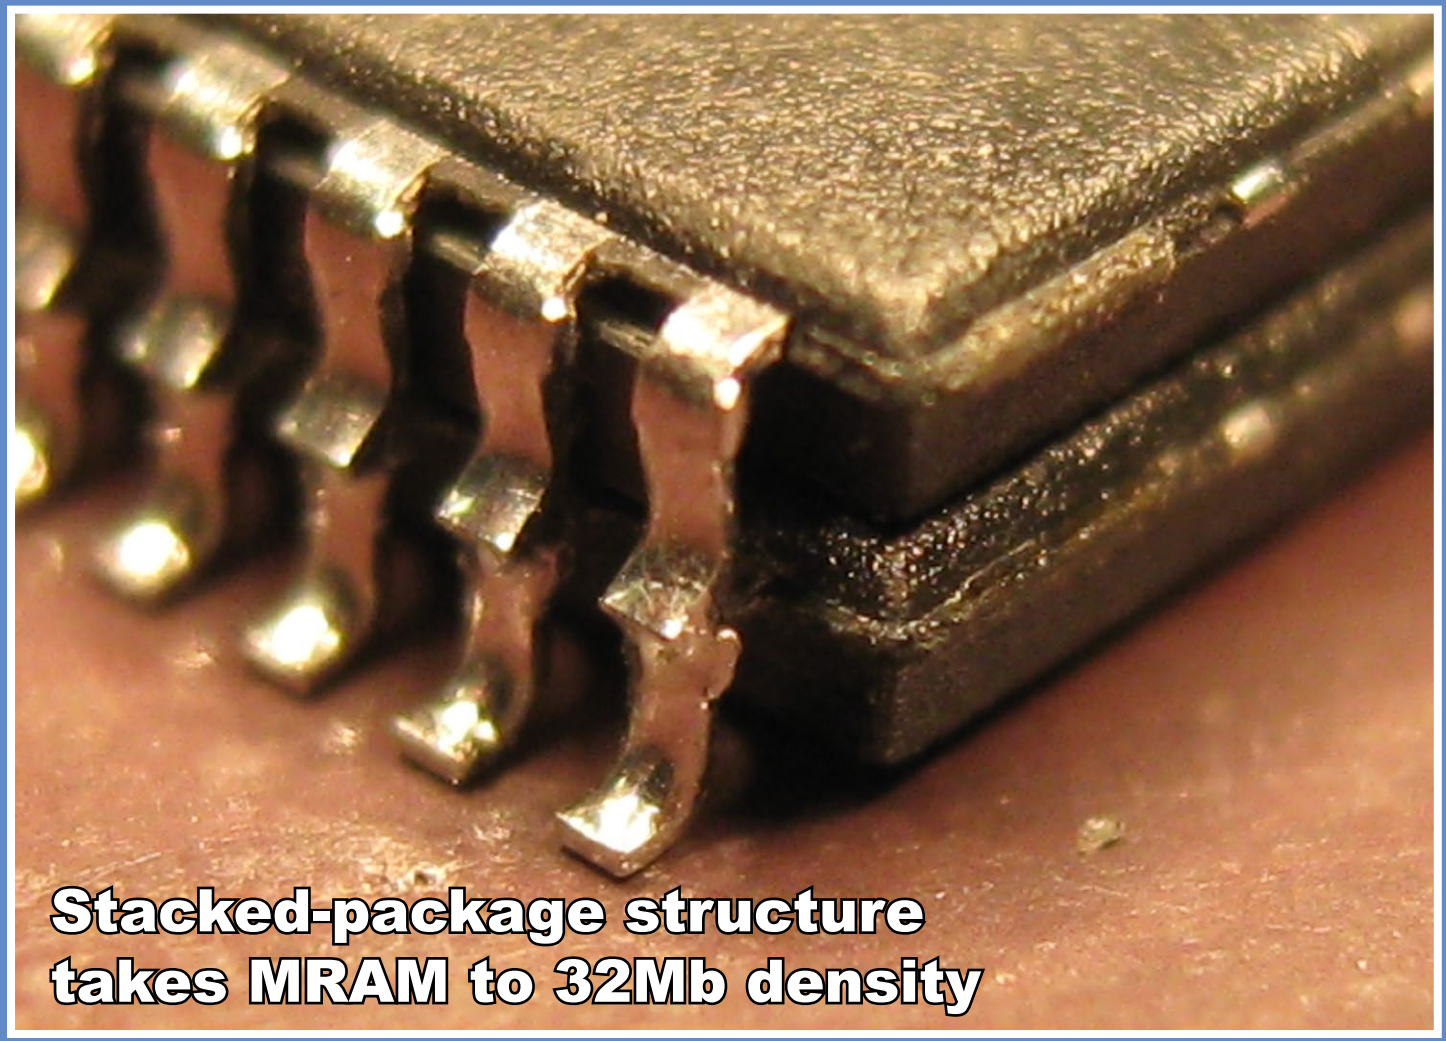

Stacked-package structure takes MRAM to 32Mb density

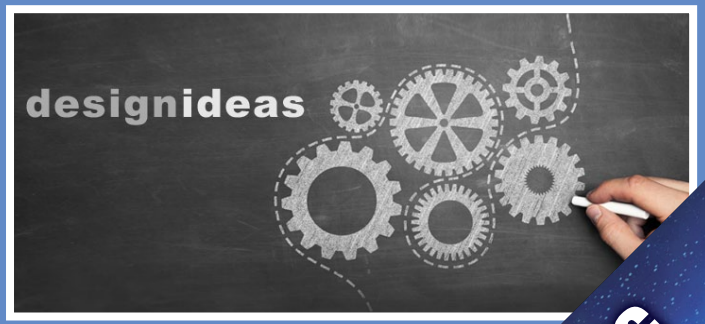

The World's Largest Selection of Electronic Components Available for Immediate Shipment!®

FREE SHIPPING
ON ORDERS OVER €65!*

GLOBAL COVERAGE

Join us at embedded world 2015

Nuremberg, Germany | 24 - 26 February 2015 | Hall 4A Booth 631

1,000,000+ PRODUCTS IN STOCK

3.9 MILLION PARTS ONLINE

650+ INDUSTRY-LEADING SUPPLIERS

100% AUTHORIZED DISTRIBUTOR

Digi-Key® CORPORATION

FIND CONTACT AND ORDERING INFORMATION FOR YOUR REGION AT

DIGIKEY.COM/EUROPE

*A shipping charge of €18.00 (£12.00) will be billed on all orders of less than €65.00 (£50.00). All orders are shipped via UPS for delivery within 1-3 days (dependent on final destination). No handling fees. All prices are in euro and British pound sterling. If excessive weight or unique circumstances require deviation from this charge, customers will be contacted prior to shipping order. Digi-Key is an authorized distributor for all supplier partners. New product added daily.

ecia MEMBER ecsn member CEDA

COVER

Highest density MRAM device built with stacked-package technique

Engineers in need of memory from the more exotic end of the spectrum now have available to them a stacked-package 32Mb MRAM-based product from e2v. The company positions the product as offering savings for the aerospace market, one of the sectors where MRAM's fast-rewrite and non-volatile performance is required. e2v says that this packaging not only offers a significantly decreased package footprint size, but also makes it the highest density MRAM device in the industry.

The EV5A16B is a stacked solution, with two 16Mb MRAM devices from Everspin Technologies and is offered in a 54-pin, stacked-package TSOP. Available in commercial (0°C to 70°C) and industrial (-40°C to 85°C) temperature ranges, the MRAM technology offers SRAM-compatible, 35-nsec read/write timing with data retention and endurance and can be used with microprocessors, DSP, storage systems, instruments, and FPGAs. e2v adds that the product addresses industry requirements for a high-density memory, available in a small footprint with the ability to increase system performance improve the size, weight and power parameters.

FEATUREARTICLES

17 **Comparing Linux, RTOS and Bare Metal for optimised real-time performance**
by Chee Nouk Phoon, Altera

21 **The basics of automotive cluster device architectures and applications**
by Deepak Mahajan, Vikas Agarwal & Arjun Pal Chowdhury, Freescale Semiconductor

25 **Does your op amp oscillate?**
by Barry Harvey, Linear Technology

27 **The fully Digital radio transmitter: Is it real?**
by Steve Taranovich, EDN

ONLINE THIS MONTH

Software tool fixes causality violations: and, De-embedding exposed
by Eric Bogatin, for EDN

Demonstrating ASIC IP performance and quality demands an FPGA-neutral design flow
by Ali Osman Örs and Daniel Reader, Cognivue

EDN's columns

4 **EDN.comment**
Yet more IoT reflections

6 **Pulse**
UWB location to 10cm: kit shows how; Cadence hits back in finFET P&R; A round of programmable parts and software from Xilinx; RS cuts 3D print prices; Multistandard IoT chip

19 **Baker's Best**
Simulating the front-end of your ADC
by Bonnie Baker, Texas Instruments

23 **Test Memories at-speed with a slow clock**
by Martin Keim, Mentor Graphics

29 **Analog Tips**
Low-power, low-cost, differential amplifier with high CMRR
by Chau Tran and Jordyn Rombola, Analog Devices

41 **7 tips for creating a reliable embedded system**
by Jacob Beningo

36 **Product Roundup**
Supercaps for energy harvesting; USB vector network analyser; Nano-amp wake-up timer; run-from-flash Renesas MCU

30 **Design Ideas**

31 Transient load gives power systems a workout

33 Monitor a differential pair for valid & fault conditions

34 NRZ to AMI converter uses single supply

YET MORE IOT REFLECTIONS

If you attended this year's Embedded World Event in Nuremberg, and you have any interest in internet-connected devices, the industry had one key message for you. That is, there were many messages you could have taken away from the conference and exhibition but one above all others. The industry (metaphorically speaking) came up close and shouted at you; it said; "You might not realise it, but you have a security problem."

To which the majority of engineers would likely have replied, "yes, we had gathered that." The realisation that the – there is no avoiding using the words – Internet of Things presents a minefield of security vulnerabilities is by now well established in most engineers' mindset. The industry, having assembled a wide variety of "components" that go to make up secure systems, is now moving towards offering more integrated solutions. It is simple enough to find the functional blocks that you might use – AES encryption is now commonplace on the most modest of microcontrollers, for example. It is worth reflecting on that point: a function that in fairly recent memory occupied a rack or at least a whole card-slot, now sits in the corner of an already small silicon die.

A keynote address at Embedded world was given by Silicon Labs' CEO Tyson Tuttle, who also identified privacy and security as key aspects of the development of the IoT as it becomes a mature market. His overall theme was that of an engineer's perspective on the technology underpinning the next steps in that market development: and his overall conclusion, that he sees no technology barriers in the immediate, or even medium-term, outlook that will impede progress. Tuttle gave his audience a "short list" of key trends to track in the short term, over the coming year. Look, he says, for a further dramatic fall in energy consumed to carry out key connected-device functions; and correspondingly, for connectivity standards that embody low-power approaches to gain ground fast. IoT SoCs – chip-scale devices with high levels of functionality – are here now, he says, and are evolving fast. And, as an aside; expect most network connections to move to IP in the near future – the benefits outweigh any extra engineering overhead, Tuttle believes.

On that note, and with the engineering perspective, Tuttle radiates confidence that the familiar Moore's Law graph will serve the IoT phenomenon well. The notion that smaller-geometry processes are only suitable for digital designs is dismissed; "RF only gets better," he says, as we progress to 40 and 28-nm processes, and there are also mixed-signal gains. How does this relate to the security issue? If, as Tuttle suggests, we have embarked on a familiar path to higher and higher levels of integration in the connected-device sector, then it seems reasonable to suppose that the security functions will follow the same path, and we will be less concerned with obtaining the individual elements of a secure solution. There will be less need to construct for ourselves a seamless and impenetrable path for data, for each individual project. Security, in Tyson Tuttle's words needs to be, "As built-in as possible."

The problem does not go away, however; the point at which we need to worry about the integrity of the data collected and processed by our connected-device networks shifts from protecting it along its path, to questions of what is done with it and who has access to it once the collection has happened and the data-mining has begun. Simply at the level of the monitoring and control possibilities that are offered for our own homes, or our health, the data takes on a different character according to the use to which it is put. You can monitor your fitness, improve energy control, save costs, increase convenience, depending on occupancy and activity levels... the propositions are by now well-known. But from an alternative perspective, that data also monitors what you are doing every moment of your day, and could become an unacceptable intrusion on your privacy. Data, as many users of social networking services are finding, is forever. At an engineering level, we have have mostly been concerned with ensuring that data is protected as it is gathered and processed; but the IoT is (or will be) a social phenomenon as much as a technological one, and systems will need to evolve in that context, and for the long term.

THINK™

ANTI-VIBRATION GROMMETS

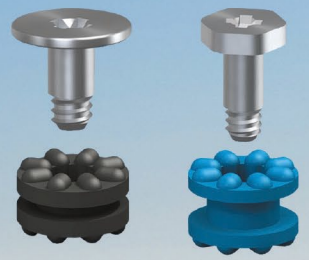

- Designed for noise suppression in office or lab equipment
- Isolates sensitive instrumentation from unpredictable vibration and shock
- Provides the flex needed to dampen unwanted noise and vibration
- Shoulder screws specially designed to provide correct spacing

IT'S WHAT'S ON THE INSIDE THAT COUNTS

KEYSTONE®
ELECTRONICS CORP.

European Headquarters: www.keyelco.com • 33 (1) 46 36 82 49 • 33 (1) 46 36 81 57

pulse

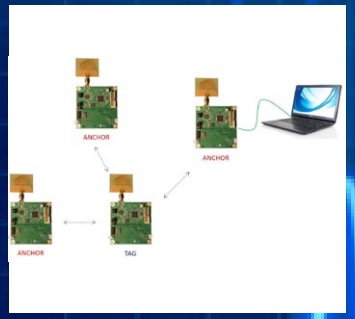
ANCHOR
ANCHOR
ANCHOR
TAG

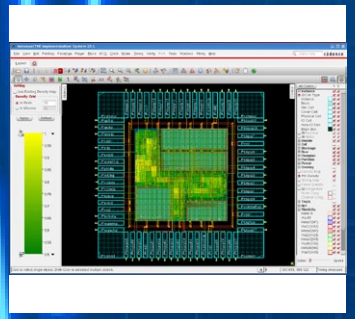

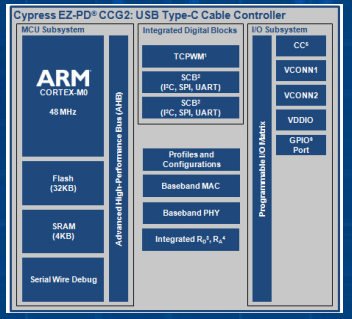
Cypress EZ-PD® CCG2: USB Type-C Cable Controller
MCU Subsystem
ARM
CORTEX-M0
48 MHz
Flash
(32KB)
SRAM
(4KB)
Serial Wire Debug
Advanced High-Performance Bus (AHB)
Integrated Digital Blocks
TCPWM¹
SCB²
(I²C, SPI, UART)
SCB¹
(I²C, SPI, UART)
Profiles and Configurations
Baseband MAC
Baseband PHY
Integrated R_X³, R_X⁴
I/O Subsystem
CC⁴
VCONN1
VCONN2
VDDIO
GPIO⁴
Port
Programmable I/O Matrix

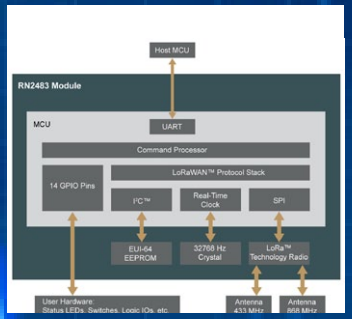
Host MCU
RN2483 Module
MCU
UART
Command Processor
14 GPIO Pins
LoRaWAN™ Protocol Stack
I²C™
Real-Time Clock
SPI
EUI-64 EEPROM
32768 Hz Crystal
LoRa™ Technology Radio
User Hardware
Status LEDs, Switches, Logic IOs, etc.
Antenna 433 MHz
Antenna 868 MHz

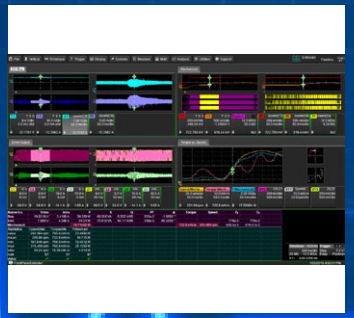


+1.123 456 8

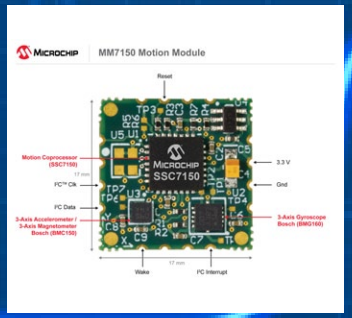
MICROCHIP
MM7150 Motion Module
MICROCHIP
SSC7150


SimpleLink™
Ultra-low Power
Wireless Microcontroller Platform
ZigBee
6LoWPAN
Sub-1GHz
RF4CE
Bluetooth
TEXAS INSTRUMENTS

Two-way ranging, real-time location evaluation kit uses UWB

Fabless semiconductor company DecaWave (Dublin, Eire) has created an evaluation kit, TREK1000, based on its precise indoor location and communication CMOS chip, the DW1000. TREK1000 enables users to evaluate the performance of DecaWave's DW1000 Ultra-Wideband (UWB) integrated circuit in different RTLS topologies. The kit will allow customers to quickly evaluate the performance of DecaWave's UWB technology in multiple real-time location systems (RTLS) use cases. Based on the two-way ranging scheme, the kit will offer the possibility to test three different topologies. The kit is based on the DW1000 chip which is an IEEE802.15.4-2011 UWB compliant wireless transceiver. The chip, while allowing the location of objects to a precision of 10 cm, is also capable of high data-rate communications, up to 6.8 Mb/sec. This combined capability makes it a perfect fit to answer the "What, When, Where" questions essential to the deployment of the Internet of Things, its makers say.

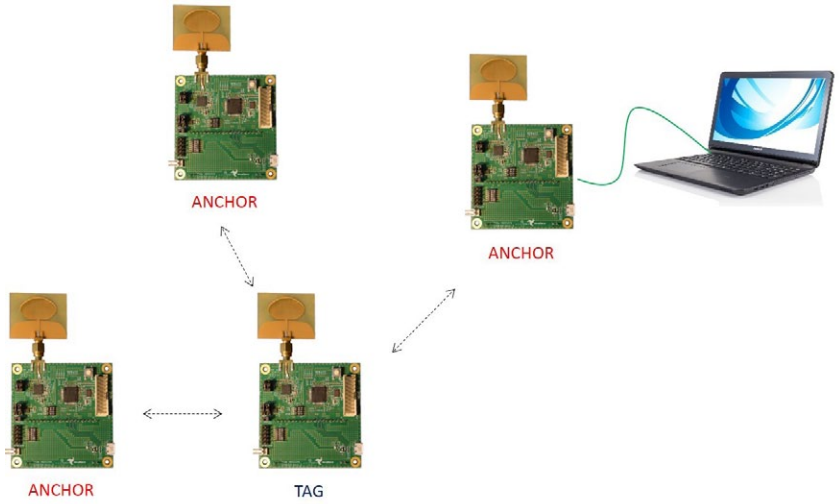

DecaWave's Mickael Viot explains that the DW1000 is quite unlike a number of UWB proposals that emerged several years ago (and that failed to take off), "Those were aimed at very high bandwidth over a very short distance to transfer large files between devices: they used OFDM whereas the DW1000 is an impulse-radio-architecture chip that makes use of the characteristics of the RF signal to maximise [distance] accuracy and reliability." It has in common with other UWB offerings the characteristic that the signal is a very low level – essentially, in the noise – and use of a coherent receiver enables reception of a useful amount of data in addition to its ranging performance. The chip is built, despite its operation over 3.5 to 6.5 GHz bands, in standard 90-nm CMOS, which enables DecaWave to be "open to" IP licensing in addition to its fabless chip supply model. Due to its communications range of up to 290m in line-of-sight and up to 35m in non-line-of-sight, it also reduces the system cost by reducing the need for infrastructure.

Complete article, here

Cadence aims to recapture share of digital chip P&R with "Innovus"

Cadence has introduced a major revision of its implementation – essentially, place-and-route – tools for large-chip design and for the finFET (16/14/10-nm) era. Among its claims for the software are significant – 10-20% - gains in power, performance and area; and massively-parallel computing assisting an up to ten-times gain in turnround time.

A Cadence spokesman openly acknowledges that in recent years, "We have been known for state-of-the-art in analogue and mixed-signal... we have not done so well in digital [chip design]" Innovus is the result of a long collaboration with foundries to produce an implementation package for next-generation technologies – but the company adds, "It's not just for finFET; the benefits will be experienced with all geometries. One

European early adopter has seen 'dramatic' improvements at 28 nm."
Innovus is "Driven by a massively parallel architecture with breakthrough optimisation technologies," - the company says that while the ability to use parallel compute resources comes from its revised approach, the speed-up is not only due to hosting on parallel compute resources; a large amount of the gain is also from the new algorithms involved.

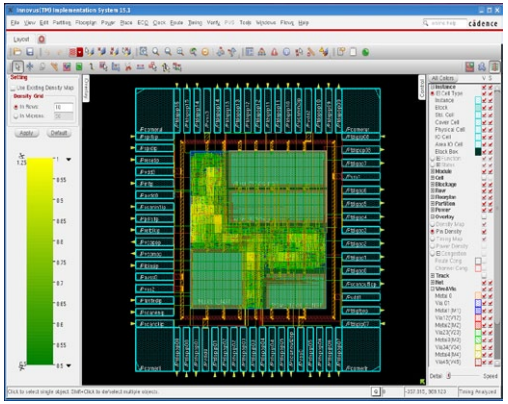

The Innovus Implementation System was designed with several key capabilities to help physical design engineers target a set power/area budget or realise maximum power/area savings while optimising for a set target frequency. Cadence's description of these upgrades lists;
- A new GigaPlace solver-based placement technology is slack driven and topology-/pin access/coloraware, enabling optimal pipeline placement, wire length, utilisation and PPA, and providing the best starting point for optimisation
- Advanced timing- and power-driven optimisation is multi-threaded and layer-aware, reducing dynamic and leakage power with optimal performance
- Concurrent clock and datapath optimisation includes automated hybrid H-tree generation, enhancing cross-corner variability and driving maximum performance with reduced power
- Next-generation slack-driven routing with track-aware timing optimisation that tackles signal integrity early on and improves post-route correlation
- Full-flow multi-objective technology enables concurrent electrical and physical optimisation to avoid local optima, resulting in the most globally optimal PPA.

Complete article, here

An "ASSP-like" S/W-driven programming environment for Xilinx' Zynq

Xilinx has announced its SDSoC development environment for its All-Programmable SoCs and MPSoCs, aiming to make use of the devices easier and more accessible while maximising their performance and capabilities. In terms of its development environments, this – the SDSoC Development Environment for All Programmable SoCs and MPSoCs – is the third "targeted" member of the company's SDx family of development environments, the previous two being centred on software-defined networks and on implementing acceleration.

In the hardware dimension, Xilinx recently extended the Zynq family of devices from "SoC" to an uprated "MPSoC" generation, adding multi-core ARM processors on-chip, among other increments. (See story here.)

With this development environment introduction, Xilinx says it has provided the assistance to designers that they need to fully exploit the SoC devices. Within it, you can take a "whole system" approach, in which the tools will assist with partitioning your design effectively, profiling the task, identifying bottlenecks, creating accelerated function blocks and migrating them to the programmable fabric so that the tasks most appropriate to software run on the embedded cores, and that best handled in programmable logic, is located there. Also included is a new measure of automation of interconnecting those functional blocks once you have determined their optimum partitioning. Automation extends to "instrumenting" the code in the system to carry out profiling, and to verify performance once development has progressed.

Xilinx takes chips to 16nm

In the update to its silicon plans, the programmable-logic maker has announced new product ranges and updates in a number of areas; forthcoming high-end product in 16-nm finFET processes; new processor options in its system-on-chip (Zynq) offering; and larger embedded memories in some of its devices. There is a virtual blizzard of new and modified product names and brand identities to penetrate; the high-end FPGAs become UltraScale+; Zynq gains multi-core hard processors ("All-programmable MPSoC"); embedded memory options are UltraRAM; and a revised interconnect strategy is SmartConnect. This release predominantly gives details of device plans ahead of availability, setting out product plans for forthcoming quarters to enable designs to begin ahead of device delivery.

A fix on that pipeline is that Xilinx says that some of its 20nm devices are now starting to move to full production: however, some of the performance (particularly, performance-per-Watt) comparisons are made from the current 28-nm product direct to the planned 16-nm step. One other piece of branding of note is that the company refers to "3D-on-3D". The first 3D alludes to the use of 3D transistors (that is, non-planar, or finFET); the second to its multi-die fabrication of its largest devices. This continues to use active devices side-by-side on a passive silicon interposer, often characterised as "2½D".

The announcement is of the 16nm UltraScale+ family of FPGAs, 3D ICs and MPSoCs, combining new memory, 3D-on-3D and multi-processing SoC (MPSoC) technologies. The UltraScale+ family also includes a new interconnect optimisation technology, SmartConnect. These devices extend Xilinx's UltraScale portfolio - now spanning 20nm and 16nm FPGA, SoC and 3D IC devices - and gain a significant boost in performance/watt from TSMC's 16FF+ FinFET 3D transistors.

Complete article, here

2 GHz signal analyser bandwidth is 'unique' says Rohde & Schwarz

Developers need large bandwidths to analyse wideband signals like those of the new IEEE 802.11ad WLAN standard, 5G next generation mobile communications signals, and radar chirp signals. The R&S FSW is the first signal and spectrum analyser offering 2 GHz analysis bandwidth. Analysis bandwidth of the FSW high end signal and spectrum analyser can now be extended to 2 GHz with the FSW-B2000 hardware option. This test solution enables R&D users to demodulate extremely wideband signals and analyse them in detail. R&S asserts that there is no other instrument on the market that combines such a large analysis bandwidth with a frequency range up to 67 GHz. To analyse a signal with such a large bandwidth, the R&S FSW downconverts it to an intermediate frequency, which is then wideband digitised by an R&S RTO oscilloscope. The R&S FSW equalises this digital signal and adjusts the sampling rate. The entire signal path including the oscilloscope is calibrated. Users control the setup via the analyser, which also displays measurement results. Measurement applications such as vector signal analysis or pulse and transient measurements make it possible to analyse the signal in detail.

Also added to R&S' offering for embedded system analysis is simultaneous time domain and spectrum analysis on 1 GHz oscilloscopes; Rohde & Schwarz' RTM bench oscilloscopes can now simultaneously analyse analogue, digital, protocol-based and RF signals. This spectrum analysis and spectrogram option can also be used in the new 1 GHz model. The RTM-K18 spectrum analysis and spectrogram option for the RTM oscilloscope family, the company says, makes RTM the only oscilloscope in its class that can analyse the time domain while simultaneously analysing the spectrum, logic and serial protocol. Interactions such as those that occur in electronic devices with RF components are quickly analysed in a single measurement. Time and spectrum analyses can be configured completely independently of one another. This means that users can simultaneously analyse signal details that differ in time and frequency, with the optimum settings for each.

Complete article, here

USB Type-C controller in chip-scale package for active USB 3.1 cables/adapters

Cypress Semiconductor has applied the combination of configurable logic, I/O driving capability and mixed-signal circuitry it uses in its programmable PsoC devices, to produce an ASSP for the rapidly-developing USB 3.1 interconnect market. The EZ-PD CCG2 Type-C controller Is Cypress' second-generation chip in this area. Now sampling is the small footprint, integrated USB Type-C cable controller solution with power delivery (PD). In USB 3, cables and interconnects are active, and on every connection event there is a negotiation to determine the upstream/downstream status of the connected device, and the power levels that are requested and delivered (or deliverable). "Type-C" designates the new, reversible miniature connector that handles the higher power and data rates of USB 3.1. Optimised for 2.4-mm thin USB Type-C cable connectors, the programmable EZ-PD CCG2 controller is fully capable of supporting any USB Type-C Downstream Facing Port (DFP) or Upstream Facing Port (UFP) applications. EZ-PD CCG2 is available in a 3.3 mm² wafer level chip scale package (WLCSP) and is claimed as the first programmable solution to fully integrate both the Type-C transceiver and termination resistors needed for Type-C communication.

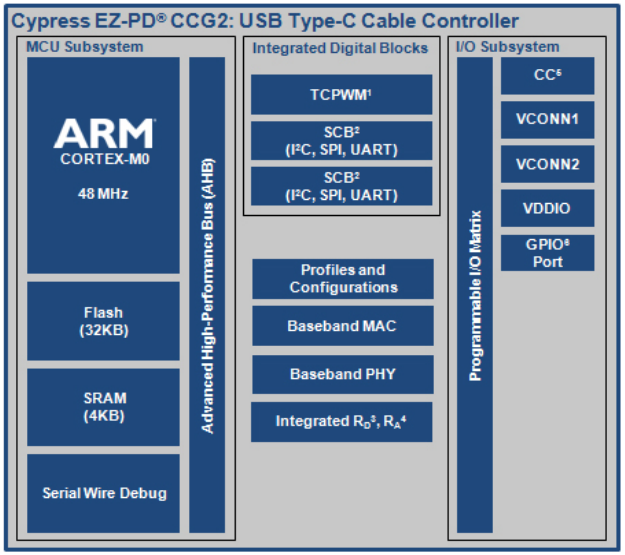

The USB Type-C standard enabes slim industrial designs, easy-to-use connectors and cables, the ability to transmit multiple protocols, and 100W PD—a significant improvement over the previous 7.5W standard. However, the Type-C standard requires an Electronically Marked Cable Assembly (EMCA) that can report the characteristics a cable supports, such as current carrying capability, protocols supported, and vendor identification. A standard MCU may be used for the EMCA implementation, but it makes the design complex by requiring multiple external ICs and passive components. Cypress says its EZ-PD CCG2 Type-C controllers solve this challenge with integration and bill-of-materials savings with only five or fewer external components required.

Complete article, here

Primary-regulated, “zero standby power”, high-power AC/DC PSU

Using the definition of zero-standby power as in the International Electrotechnical Commission (IEC) 62301 clause 4.5, Texas Instruments has assembled a power supply design with smart load monitor, and wake-up feature that enables rapid step response with smallest output capacitors for 75-W flyback converters. Presented as the first zero standby power controller chipset with the lowest standby power consumption for AC/DC flyback power supplies up to 75W, the UCC28730 primary-side regulation (PSR) flyback controller with 700-V startup switch and UCC24650 wake-up monitor supports 5- to 24-V output voltages and helps designers create smaller, more efficient power supplies.

A major design limitation with existing primary-side regulation power supplies, TI adds, is poor transient response performance while trying to maintain low standby power. The UCC28730 and UCC24650 chipset achieves zero standby power consumption while achieving the best transient response performance possible. Lowest system-standby power consumption enables zero standby power consumption without any additional components in standard flyback designs up to 20W and is extendable up to 75W using the intelligent sleep function for adaptive relay control.

Complete article, here

Open source, 1U rack mount internet appliance, based on Raspberry Pi

Californian company EarthLCD recently introduced what it says is the only-available 10-in. wide by 1-in. high TFT LCD panel, which it uses to give single-unit-height rack modules a versatile display; now the company has put together an open-source design for a rack-mount internet appliance based around the Raspberry Pi. The unit is available as open-source project files – or as a complete module. In 2014, EarthLCD proposed the idea of a 1U rack mount open source appliance that would enclose its unique 10:1 aspect ratio TFT LCD. In order to get the LCD rapidly designed in, EarthLCD took the step of creating not just a reference design, but completing an open-source design which includes packaging, software and firmware to build a 1U rack mount appliance.

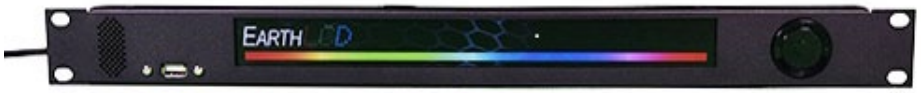

The Pi-Raq has a powder-coated sheet metal enclosure holding a Raspberry Pi B+ with an EarthLCD 10 x 1 LCD adapter board, the EarthLCD-10.4-1024100 TFT colour display, and a user interface control with a jog wheel. The SolidWorks (3D CAD) files are included for this metal enclosure along with the schematic and [PCB] Gerber for the LCD driver. With the Pi-Raq and the existing EarthLCD 10 x 1 display, rack mount equipment, such as power controllers, network hubs and routers, can now have an easily programmable front panel with colour TFT LCD.

Complete article, here

LoRa wireless module for IoT designs, from Microchip

Microchip Technology's first module for ultra long-range and low-power network standard is a stack-on-board package; the RN2483 module is intended to simplify access to the 10-mile range and 10-year battery life of LoRa technology wireless networks.

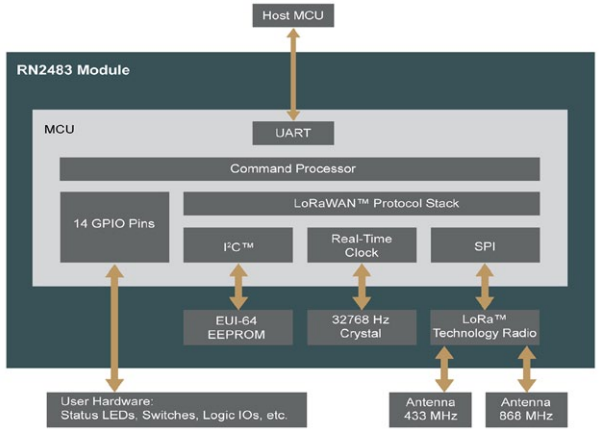

The module is the first in a planned series for the LoRa technology low-data-rate wireless networking standard, which enables Internet of Things (IoT) and Machine-to-Machine (M2M) wireless communication with a range of more than 10 miles (suburban), a battery life of greater than 10 years, and the ability to connect millions of wireless sensor nodes to LoRa technology gateways. The 433/868 MHz RN2483 is a European R&TTE Directive Assessed Radio Module, accelerating development time while reducing development costs. It combines a small module form factor of 17.8 x 26.3 x 3 mm with 14 GPIOs, providing the flexibility to connect and control a large number of sensors and actuators while taking up very little space.

LoRa technology has several advantages over other wireless systems. It uses a spread-spectrum base modulation that is capable of demodulation with a 20 dB below noise level. This enables high sensitivity with robust network links, improves network efficiency and eliminates interference.

Complete article, here

RS stirs 3D printer market with launch of own-brand, budget-price model

Presented as a professional-quality 3D printer for rapid prototyping at highly affordable price point, the RS IdeaWerk 3D printer targets a wide range of users including electronics engineers, enthusiasts and students.

Distributor RS Components (RS) says its first RS brand 3D printer offers professional-quality specifications and exceptional print output at a price of only €737/£599, making it approximately 30% cheaper than other 3D printers in its class.

The IdeaWerk employs FDM (Fused Deposition Modelling) technology and offers high-level specifications including a build volume of 150 x 150 x 140 mm and a minimum layer thickness resolution of 0.18 mm. The single-head system accepts 1.75 mm-diameter PLA filament materials with many different colours, also available from RS.

The unit is, adds RS, lightweight enough to be portable (7.5kg), and has dimensions of 211(L) x 403(W) x 298 mm(H). It can be used in either online or offline mode, i.e. with or without a PC with connectivity via SD Card or USB. The printer is compatible with Mac OS and Windows OSs including XP, Vista, 7 and 8/8.1, with no additional software or accessories necessary.

Complete article, here

Programmable analogue chips as I/O to MCU & FPGA-based systems

At the recent Embedded World exhibition, Maxim Integrated was reporting a broad range of interest in its MAX11300 configurable analogue device, also known as "PIXI". We first reported the device in EDN Europe June 2014, and Maxim is currently promoting it together with its dedicated programming tools.

Over many years there have been numerous attempts, with varying levels of success, to build programmable analogue parts. Some of the problems they have encountered have included limited (parametric) performance, functional blocks restricted to basic components (tiled arrays of op-amps and passive components) and lack of support in programming. Maxim says it has addressed all of these to make a fully-programmable, mixed-signal I/O chip.

You get 20 configurable, high-voltage, bipolar ports, each of which can be an ADC analogue input, a DAC analogue output, a general-purpose input port, a general-purpose output port, or an analogue switch terminal. The part combines a 12-bit, multichannel ADC and a 12-bit, multichannel, buffered DAC—all in a single, programmable IC.

The company reports that engineers are using the part to create customised I/O circuitry where no catalogue part exists that exactly matches the requirement; this can replace multiple discrete devices and passives, and save board area. Where a product line has a variety of input and output options, a single part and a single PCB layout can cover a range of product part numbers (say, the option of a 0 – 10V or a 4 – 20 mA output on a controller). A further option is to have that level of configurability in a single product, under software control.

Complete article, here

Power analyser provides complete analysis for 3-phase motor drives

Teledyne LeCroy's MDA800 Series of Motor Drive Analysers (MDAs) combine three-phase power analyser static (steady-state) calculations, dynamic three-phase power and mechanical motor analysis capabilities, and high bandwidth (1 GHz) embedded control system debug into a single instrument.

The Motor Drive Analysers (MDAs) are based on the HDO8000 oscilloscope platform and are standard with eight input channels (16 digital channels optional) with 12-bit resolution, 2.5 Gsample/sec sample rate, up to 1 GHz bandwidth and up to 250 Mpts/ch acquisition memory. A complete set of serial trigger/decode and analysis software options, and a wide variety of voltage and current probes are available to use with the MDAs.

Positioning the instrument, LeCroy notes that a drive system combines three-phase power electronics, motor/mechanical analogue and digital sensors, and embedded controls, with a complex variety of analogue, digital, serial data and pulse-width modulated (PWM) signals. Conventional 8-bit oscilloscopes are appropriate to capture higher-speed embedded control system or power transistor activities: the [conventional] power analyser is a focussed tool for measuring input/output "black-box" drive power and efficiencies, but provides limited or no waveform capture for embedded control or drive system debug.

The MDA permits waveform captures from a drive's three-phase power section, individual power transistors, and embedded control system, and performs coincident three-phase power analysis of the power section waveforms in one high-performance instrument, enabling debug and analysis of all aspects of the complete motor drive.

Complete article, here

6½ & 7½ digit DMMs from Keysight visualise measured data

Two "Truevolt" Series digital multimeters (DMMs) – the Keysight 34465A DMM (6½ digit) and the Keysight 34470A DMM (7½ digit) – help engineers visualise measurement data in multiple ways. Truevolt DMMs' advanced graphical capabilities, such as trend and histogram charts, help engineers achieve greater insights faster. Both models offer three acquisition modes: continuous running for typical measurements, data logging for easier trend analysis, and a digitising mode for capturing transients. The 34465A DMM is a new, higher-performance 6½ digit class of DMM that Keysight presents as providing higher speed, better accuracy and more memory. The 7½ digit 34470A DMM is a new product category of DMMs for Keysight and provides even greater resolution and accuracy, a requirement for today's most challenging devices. Both DMMs offer deeper memory for data storage than previous generation DMMs. They also provide the ability to measure very low current, 1 µA range with pico second resolution, for measurements on very low power devices; plus additional measurement and math functions, extended measurement ranges and advanced triggering.

Complete article, here

Microchip goes modular for motion measurement

Microchip Technology's 17 x 17 mm, MM7150 Motion Module combines a Microchip motion co-processor IC with MEMS sensors by Bosch. The module integrates a 6-axis digital compass, 3-axis gyroscope and a magnetometer. The Microchip silicon comprises the SSC7150 motion co-processor, which comes pre-programmed with the required sensor fusion algorithms. With an I²C connection to most MCUs/MPUs, embedded/IoT applications can make use of the module's advanced motion and position information, which is available as "raw" or processed data outputs.

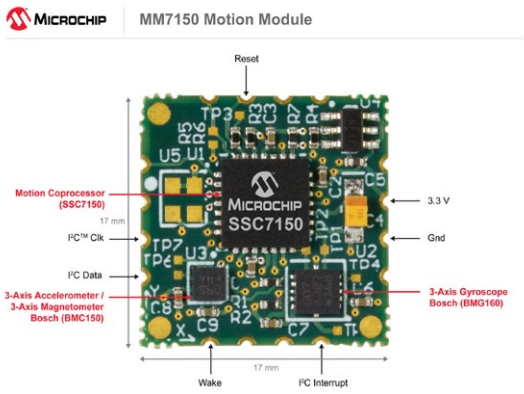

The software filters, compensates and combines raw sensor data to provide accurate position and orientation information. The small form factor module is self-calibrating during operation utilising data from the pre-populated sensors: the Bosch BMC150 6-axis digital compass; and the BMG160 3-axis gyroscope.

The motivation for this development arose from the prevalence of motion-oriented applications that have arisen since motion sensing became prevalent in smartphones; and from Microsoft's specification of a standard format for motion data to be imported into Windows. Not all application developers who might want to include motion/position may have access to the resources to develop the algorithmic aspects, Microchip believes: hence the module concept. Microchip can sell the SSC7150 IC alone for specific integration projects, but notes that with the module, all the mapping to correct for position and orientation variations is done for you.

Complete article, here

Embedded HMI package moves to higher resolution and image quality

High-resolution EVE (embedded video engine) products from FTDI Chip can now deliver faster upload speeds, greater colour depth and improved video playback for next generation HMIs; the use of 7-in. graphics panels in industrial and consumer product designs is a particular target.

The EVE range of products for human machine interface (HMI) implementation gains four new ICs - the FT81xQ series. These EVE devices support higher maximum screen resolutions than previously possible - increasing from the 512 x 512 pixels of the company's FT800Q/801Q offering up to 800 x 600 pixels, in order to address larger displays (7-inches and above).

The FT810Q has 18-bit RGB interfacing and resistive touch functionality, while the FT811Q has 18-bit RGB interfacing and is designed for capacitive touchscreen implementation (with provision for 5-point touch detection). These are complemented by the FT812Q and FT813Q for use respectively with resistive and capacitive touchscreens, but each with 24-bit RGB.

As a result of algorithm enhancements, the devices in the FT81xQ series have much smoother video playback. Screen rotation through 90° is easier to achieve, permitting both landscape and portrait orientations - a clear benefit when these ICs are utilised in handheld designs. Rotation is now a single command – previously, the designer had to rotate each object (EVE is an object-oriented graphics environment) individually. The memory capacity has also been increased from 256 kBytes up to 1 Mbytes and a quad SPI interface accelerates data transfer rates.

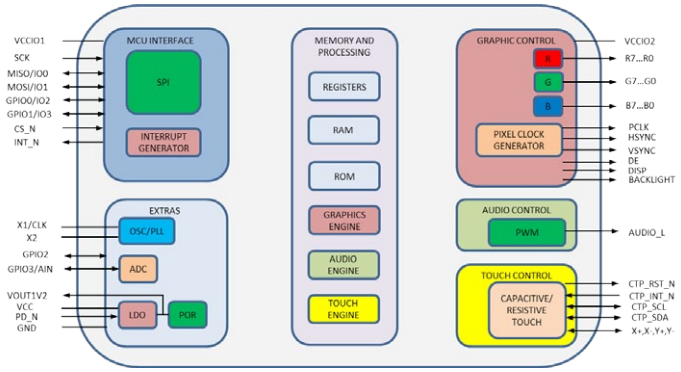

"It is now two years since the first EVE devices were introduced to the market and in that time we have seen considerable buy-in to the concept, due to the simplified HMI system architecture that EVE enables. Now with these latest additions to the offering we are giving engineers higher resolution, greater colour depth, larger memory capacity and audio input, so that they can construct HMIs that have a much higher degree of sophistication and enhance the overall user experience," explains Fred Dart, CEO and Founder of FTDI Chip.

Complete article, here

Battery-less IoT connectivity with multi-standard wireless MCU

TI has introduced the first three in a family of SimpleLink ultra-low power wireless microcontrollers that are intended to run for years on small batteries, or operate from harvested (ambient) energy; they have Bluetooth Smart, 6LoWPAN, ZigBee, sub-1 GHz and ZigBee RF4CE air interfaces. With these devices, TI says that designers have the flexibility to develop products that support multiple wireless connectivity standards using a single-chip and identical RF design. The SimpleLink ultra-low power platform supports Bluetooth low energy, ZigBee, 6LoWPAN, Sub-1 GHz, ZigBee RF4CE and proprietary modes up to 5 Mbps.

The SimpleLink ultra-low power platform is – TI asserts – the most integrated with an ARM Cortex-M3 MCU, Flash/RAM, analogue-

to-digital converter, peripherals, sensor controller and built-in robust security on chip. The platform comes with ready-to-use protocol stacks, TI RTOS, and Code Composer StudioIDE.
The first members of the SimpleLink ultra-low power wireless MCU platform are the CC2640 for Bluetooth Smart, and the CC2630 for 6LoWPAN and ZigBee. The CC2650 wireless MCU supports multiple 2.4 GHz technologies including Bluetooth Smart, 6LoWPAN, ZigBee and RF4CE. The multi-standard support helps future-proof designs and configure a chosen technology at the time of installation in the field. Additional members of the platform – the CC1310 for Sub-1 GHz operation and the CC2620 for ZigBee RF4CE – will be available later in 2015.

Complete article, here

COMPARING LINUX, RTOS AND BARE METAL FOR OPTIMISED REAL-TIME PERFORMANCE

By Chee Nouk Phoon, Altera

Real-time systems can need uncompromising hard real-time requirements where the jitter on deadline has to be within certain bounds. Others present soft real-time requirements, such as optimised energy efficiency, which will not introduce catastrophic failure but are still very important over long periods of operation.

Either way, it is important to understand the exact real-time response of a given system architecture in terms of real-time loop latency, jitter and other requirements. It may initially be assumed that a 'no-OS' or 'bare-metal' system will inherently be faster, but this is not necessarily true and running an application on a very-high-speed application processor running a rich RTOS may actually provide a better response time than a bare-metal implementation.

Requirements and challenges

Industrial control offers an exemplary case study for real-time application requirements, as it comprises both real-time and non-real-time tasks. The real-time tasks handle external interrupts – either via register polling or interrupt services – which occur in the order of tens of microseconds to respond to the interrupt, move the necessary data, do computation, and return results before the next interrupt. The jitter cannot exceed a few microseconds to ensure real-time response.

Users will often group all the real-time processing on to one particular core for more direct control and ideally to obtain higher performance. The non-real-time tasks typically include housekeeping tasks, networking and the user interface. There is little or no sharing of peripherals between processor cores, but it is necessary to share some common memory buffers for synchronisation, communication or display data. Other important requirements include ease of programming and risk minimisation, making it essential to have strong ecosystem support – in this case it is the ARM ecosystem. Also, so that a design can quickly move to more powerful processors and benefit from software innovations, it is important to have portability in hardware and software design, which often means programming above an OS abstraction.

Real-time response and jitter tolerance requirements dominate most real-time design decisions. Response time is typically expressed in terms of a real-time loop during which the system has to handle an interrupt and perform all the required computing before the arrival of the next interrupt (see Figure 1). The loop can vary from approximately 1 µsec in a software real-time system to several or tens of microseconds in real-time systems.

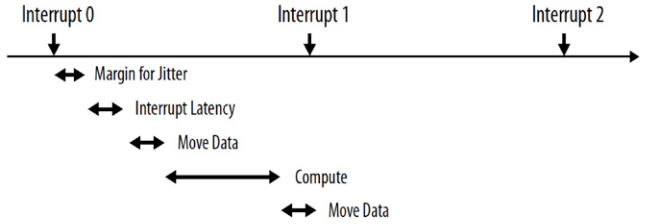

Figure 1. *Real-time loop*

Application example

An example application, running on an Altera Cyclone V SoC development board, can provide data for an objective evaluation of different real-time OS configurations. The Cyclone V SoC integrates a high-performance application processor with FPGA fabric (see Figure 2).

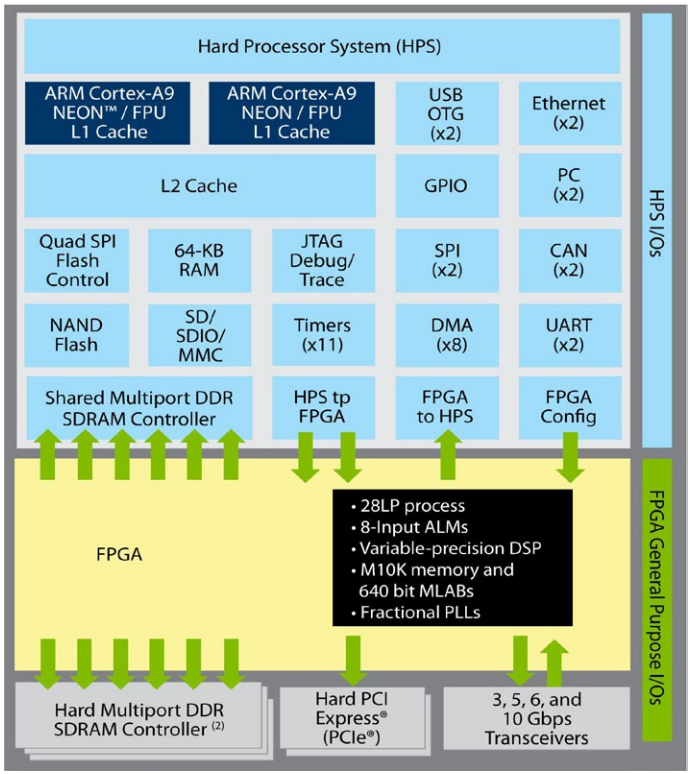

Figure 2. *Altera Cyclone V SoC*

The dual-core ARM Cortex-A9 processor present in the SoC hard processor system (HPS) is tightly coupled using the ARM MPCore technology in a classical symmetric multiprocessing (SMP) configuration. Mature and proven software solutions are available from many ARM ecosystem software providers to enable asymmetric multiprocessing (AMP). The example application system uses both the dual-A9 cluster and the FPGA and there is also a small DMA 'DataMover' design that runs on the FPGA, working in tandem with the A9 cluster to move the data to and from the FPGA.

The system receives interrupts from the FPGA and data is sent from the FPGA to the HPS for handling, which involves a small amount of computation, after which some results are written back to the FPGA. The tasks that simulate interrupt handling, data moving, and data return are the 'real-time tasks'. The round-trip loop time and jitter are measured as an indication of the system's real-time responsiveness and interrupts are handled via polling or interrupt service routines (ISR).

Figure 3 shows the reference hardware design. In this system, external interrupts are simulated by data stored on an on-chip RAM implemented in the FPGA. The size of the data can be changed easily to study the relationship between data size and efficiency of DMA operations.

The application software was implemented in three different configurations: Linux (LTSI v3.10) and VxWorks (v6.9) individually running in SMP-core-affinity mode running on both cores of the

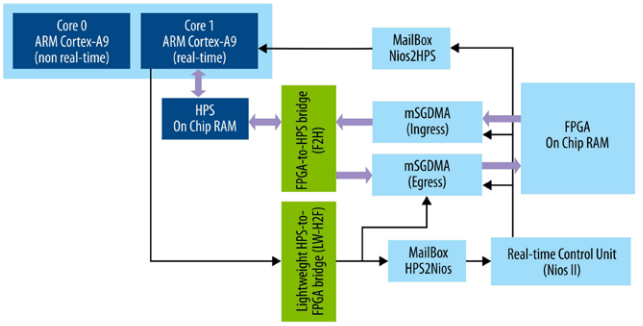

Figure 3. *Real-time application system architecture example*

dual-A9 cluster; and thirdly, bare-metal mode using hardware libraries in SoC Embedded Design Suite 14.0, running on a single core. All three implementations use Core 1 for the DataMover for data movement and interrupt handling, while the two OS implementations also use Core 0 to run a continuous Fibonacci series to simulate non-real-time tasks. Core 0 is not used in the bare-metal case to represent a 'best case' environment with no overhead or buffer management.

Read the results; Download the pdf of the complete article, which the author concludes with charts of the performance of the example system in each of its configurations.

BY BONNIE BAKER, TEXAS INSTRUMENTS

Simulating the front-end of your ADC

Successive-approximation, analogue-to-digital converters (SAR-ADCs) are straightforward, right? You attach analogue voltages to the inputs (AINP, AINN, REF), and you see an output digital code that represents the analogue input voltage with respect to the reference. At this point, you may be tempted to analyse the converter's specifications to verify that your converter operates up to datasheet standards. Not so fast! Are you sure that the converter has really received the correct analogue signals internally?

You can anticipate and troubleshoot these types of problems by using simulation tools. The determination of the ADC analogue input stage simulations rely on voltage and current accuracy. This is where the analogue SPICE macro-models come in handy. The PCB digital signal-integrity relies on timing, voltage-current levels, and parasitics. This is where the digital IBIS model comes into play. The discussion about IBIS is coming next time, but let's address the simulation environment with ADCs.

SPICE simulations for ADCs

A trial-and-error approach to sending signals into your ADC is time-consuming, and it may or may not work. If your analogue input pins are not stable at the critical times when the converter is capturing the voltage information, it will be impossible to obtain the correct output data. The first step that the SPICE model allows you to do is verify that all analogue inputs are stable so that there are no erroneous signals going into your converter.

Let's look closely at a typical serial, pseudo-differential SAR-ADC like the ADS8860 (Figure 1).

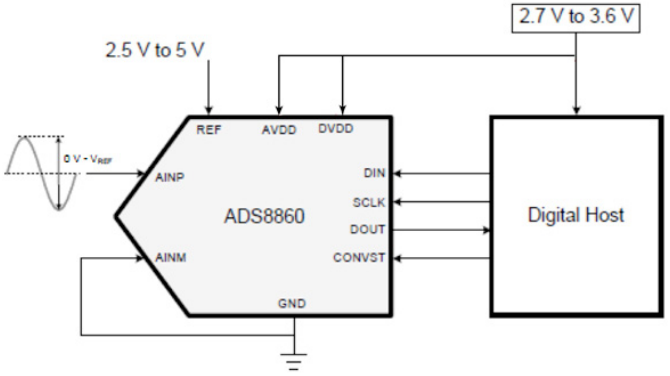

Figure 1 *The ADS8860 is a pseudo-differential input, 1 MHz, 16-bit SAR-ADC.*

This device's TINA-TI spice macro model allows you to simulate the effects of the analogue signals going into the converter. With this model and the proper driver op amp models at AINP, AINM, and REF, you can determine whether a good conversion is possible before you go to the PCB. The importance of the ADC macro-model is that it characterises the converter's input terminals accurately. The op amps driving AINP, AINN, and REF must also model their open-loop output resistance (Ro) accurately.

Let's get busy in figuring out how this macro-model works. The converter macro-model samples both positive and negative inputs individually with 55-pF sampling capacitors. The device converts the voltage difference between the two sampled values at AINP and AINN. As you view the simulation results, the model must settle to at least half a LSB at the end of the acquisition period. For this 16-bit converter half a LSB equals REF / 2^{16}.

The voltage reference pin, REF, requires a stable voltage to be present during the conversion process, or after the CONVST pin becomes a high value (Figure 2). While CONVST is low, the converter is acquiring the input signal (acquisition mode). The SAR-ADC macro-model has a 1 MHz clock and does produce the CONVST signal. The voltage reference pin must settle at the end of the bit conversion periods throughout the entire conversion time

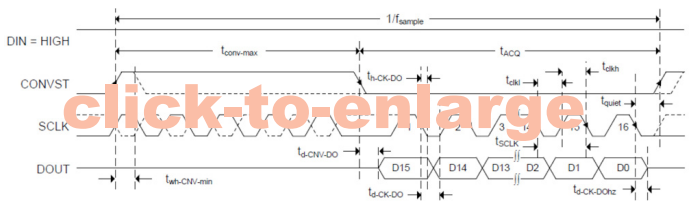

Figure 2 *In this three-wire timing diagram with three-wire operations, CONVST functions as chip select.*

The TINA-TI model for the ADS8860 in Figure 3 samples the input signals on AINP and AINN and presents the results on the model's AINPsmpl and AINMsmpl.

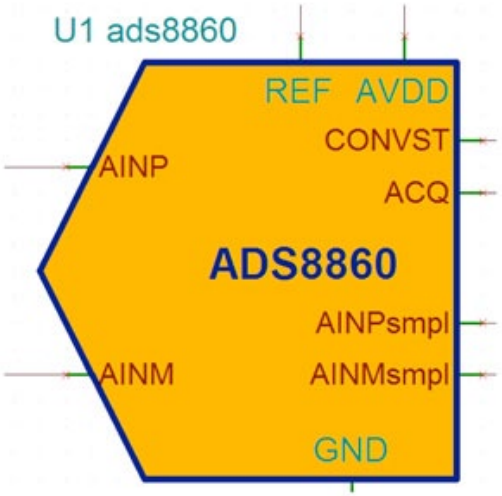

Figure 3 *Here is a TINA-TI macro-model of a SAR-ADC.*

In Figure 4, the input at AINP is equal to 3V, and the reference voltage is equal to 4.096V. As you test the accuracy of the input signal, set the ADS8860 TINA-TI circuit to sense the difference between the output of the amplifier driver, AMP_OUT_sig, and its output signal, AINPsmpl. As you look at this difference, examine the region at the end of the acquisition time or just before the CONVST pin goes high. Verify that the signal is less than half a LSB.

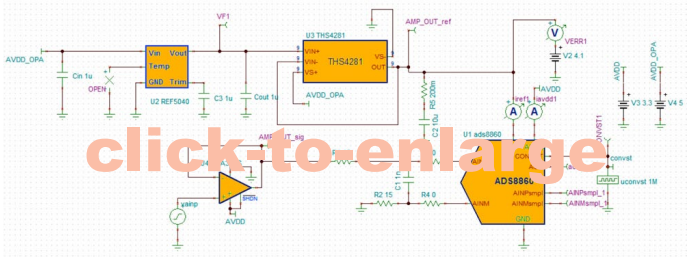

Figure 4 *This set-up is for the TINA-TI circuit to monitor the analogue and reference inputs.*

Once you have examined the analogue input function for accuracy, examine the stability of the voltage reference pin. As you measure accuracy of the REF pin, measure the difference between the voltage reference output (VF1) and the THS4281 amplifier output (AMP_OUT_ref). Be sure to remove offset errors generated by the voltage reference (REF5040) and op amp (THS4281) with the value of VERR1. While doing this, examine the voltage level just prior to the current spikes using the iref1 current meter. Reference 1 provides good in-depth information for this simulation.

Simulation is a tricky thing when it comes to SAR-ADCs. Presently there is no complete converter model that accurately models the entire device. The resource you do have is an analogue SPICE file that models the analogue input pin stability. The good fortune of having this tool available is that you have a powerful tool to address one of the most critical, difficult converter issues.

But, this is half the story. You have only simulated the analogue part of the ADC. Next time we will talk about simulating the digital input / output of your converter.

REFERENCES

1. "Using SAR ADC TINA Models: Much ado about settling," Munikoti, Harsha, Precision Hub, Texas Instruments, December 12, 2014
2. Download the ADS8860 datasheet.

THE BASICS OF AUTOMOTIVE CLUSTER DEVICE ARCHITECTURES AND APPLICATIONS

By Deepak Mahajan, Vikas Agarwal & Arjun Pal Chowdhury, Freescale Semiconductor

With the increased complexity of vehicle electronics, greater functionality requires status information to be displayed to the driver. The instrument cluster is the primary data source for the driver, delivering information about vehicle and engine status. Given system complexity, however, there is greater demand for a more user friendly, lucrative and cost-effective solution to support a wide range of automotive cluster applications. Here we will discuss various components of a cluster device that enable this support. The article classifies cluster architecture and applications into the following broad categories:

- Types of cluster devices
- Autosar application component
- Graphics application component
- Cluster security
- Device memory requirement
- Low power mode cluster architecture

Cluster device types

There are three basic cluster device platforms, namely:

Cluster Graphics Platform; Car manufacturers and tier-1 suppliers are facing an increasing need for content to be displayed to the driver requiring instrument cluster solutions that can be dynamically reconfigured to display content based on context and driver preferences while presenting essential information to ensure safe driving. A Cluster Graphics Platform is mainly used to fulfil this requirement. This device mainly consists of a Graphics processor with 2D/3D accelerator, huge on-chip graphics RAM for high bandwidth graphics data access, high-speed external memory interface such as SDR, DDR or NAND Flash Controller, and a display controller to display graphics information through a TFT, LCD display. In addition, a set of integrated peripherals such as CAN, MOST MLB, LVDS, Ethernet and USB allow a direct connection to the rest of the car network. The device is generally used for high end cluster application by OEMs.

Cluster Controller Platform; This is the cluster MCU device that mainly handles the Cluster AUTOSAR stack's functionality. This device contains a stepper motor controller and driver for gauge driving, sound generation module, PWM channels for sound and cluster LED control, various analogue sensors and communication peripherals, HMI to communicate with driver, other body devices and gateway devices or other subsystem. The device can also have segmented LCD or even dot matrix LCD displays as an alternative display solution for low-end cluster application. For high/mid range applications, most of the real time data and status captured by the device is communicated to the cluster graphics device through fast SPI or EBI interface.

Combined Cluster Platform; The graphics controller and real time application controller (Cluster MCU) is combined in this platform to make a single chip solution for both Graphics and Autosar application. This device keeps a balance between graphics performance and real time performance and is therefore used by OEMs when optimum solutions are required. Some renowned semiconductor companies offer a wide range of combined cluster platform

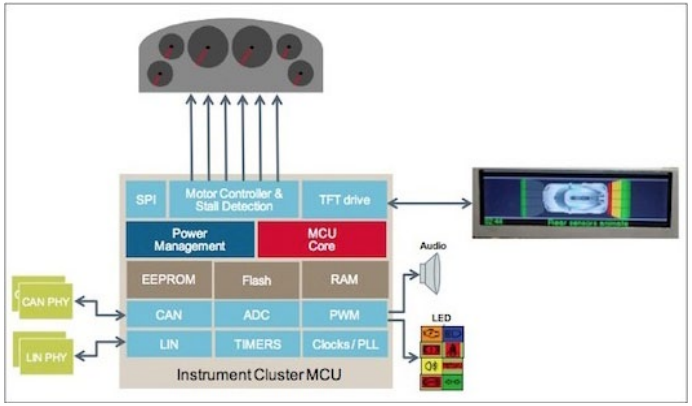

microcontrollers with multi core based solutions to support from basic to premium line instrument cluster applications.

Module use-case; AUTOSAR application components

Core: The core is part of real-time application domain of the device. It is mainly used to run AUTOSAR applications such as basic a communication driver, PWM driver, handling and servicing of various peripherals etc. The high bandwidth motor control application driver can also be controlled through this core. The device can also have a dedicated core for motor control applications, which helps to share the bandwidth of the real time processor.

Stepper Motor Control/Driver: In a cluster application a stepper motor driver is mainly used to control gauges. Real-time information of vehicle speed, fuel level, tachometer, temperature etc., is communicated to cluster devices through a gateway device. The core then processes that information and updates the motor control driver to control the stepper motor. Two motor control channels are mainly used to drive a single stepper motor.

PWM: PWM outputs are used to control the backlighting of LCD displays, gauge backlighting etc. The LED intensity of dashboard, odometer, etc., are controlled by PWM outputs. For low-end applications where few LED tell-tales are required, it is controlled by a microcontroller

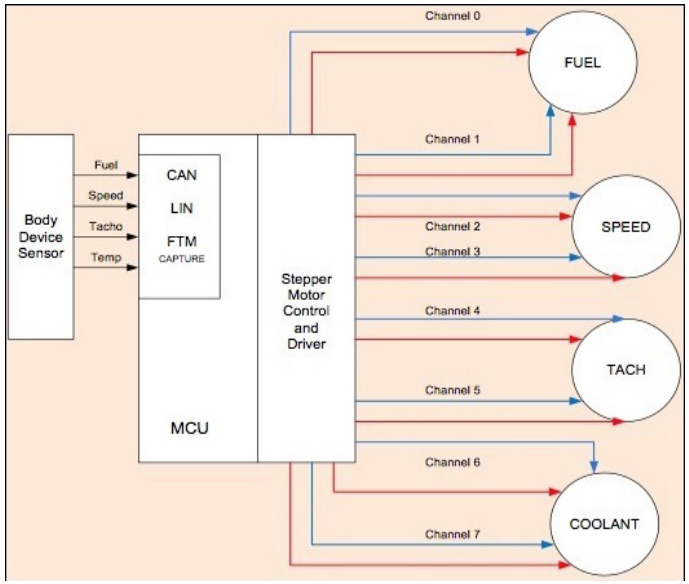

PWM output. However in mid-range or high-range applications where more LEDs are present in the system, there is typically a dedicated LED driver.

ADC and Alarm Comparator: These analogue components are mainly used to monitor the currents and voltages of various power supplies. They are also used to measure the current flow through Telltale LEDs to check that everything is working correctly. Measuring the battery level, ignition level, air temperature and coolant are common ADC functions performed by either a cluster or body device. When performed by a body device, it is communicated to a cluster by using a LIN or CAN connector.

Sound Generation Module: The device can have a sound generation module with I²S protocol support, mainly to produce a buzzer sound in a cluster to indicate the occurrence of some unwanted event. For example, a driver seat belt warning detected by a body-device sensor is communicated to a cluster. The sound generation module inside the cluster device will create a buzzing sound to alert the driver of the event. Similarly, door lock/unlock and other common events can be communicated through a sound generation module. It can also be used to play MP3 audio data with the help of a CPU that converts MP3 into PCM data for the I²S interface.

The authors continue this comprehensive overview of what it takes to assemble a next-generation automotive instrument cluster, looking at communications interfaces; graphics components; the applications processor, and associated GPU and display controller. After considering come of the detail aspects of the graphics system – including video handling – the article concludes by noting the trend to a future highly-integrated system.

TEST MEMORIES AT-SPEED WITH A SLOW CLOCK

BY MARTIN KEIM, MENTOR GRAPHICS

Recently, I was involved in supporting an MBIST (memory built-in self-test) user. Because of an error in the design, the engineer couldn't run a memory test because the high-speed system clock wasn't available. In this case, only the relatively slow test clock was available. Of course, the user was very concerned about the quality of the memory test and was even more concerned about the potential increased DPPM (defective parts per million) number of his product. Fortunately, most memory tests aren't dependent on the high-speed clock signal.

Using a slow-speed clock, the chance of detecting memory defects reduces very little, which results in [only] slightly higher DPPM levels. Whether this higher DPPM level is significant for the product depends more on the product's application than on the test. For automotive or medical products, even the slightest increase in DPPM is unacceptable, but the same DPPM increase for low-cost consumer electronics might very well be within the contractual obligations.

The ability of modern memories to self-time is at the core of the mystery. Self-timing is the ability to execute a read or write operation at a certain speed without dependency on an external clock stimulus. The time starts when a change of certain memory control input ports signal the start of a read or write operation. The time then stops when the operation is complete.

There are two important paths in the memory that determine the test: the path the data needs to take and the self-timing path. The purpose of the self-timing path is to always produce the same delay, within margins, and then trigger the sensing of the data coming out of the memory array. Together, these paths set the speed at which a system's memory operates reliably.

To be precise for the context here, synchronous, embedded SRAM (Static Random Access Memory), used in today's microelectronics, are virtually all self-timed. Figure 1 depicts a self-timed memory. The blocks and gates shaded grey are the key components of the self-timing feature. The delay through the Model Row and Column Decode logic determines how long the write drivers are turned on during a write operation and when the latches should capture the output of the sense amplifiers during a read operation, after the occurrence of a rising clock edge. Once the operation is complete, the address

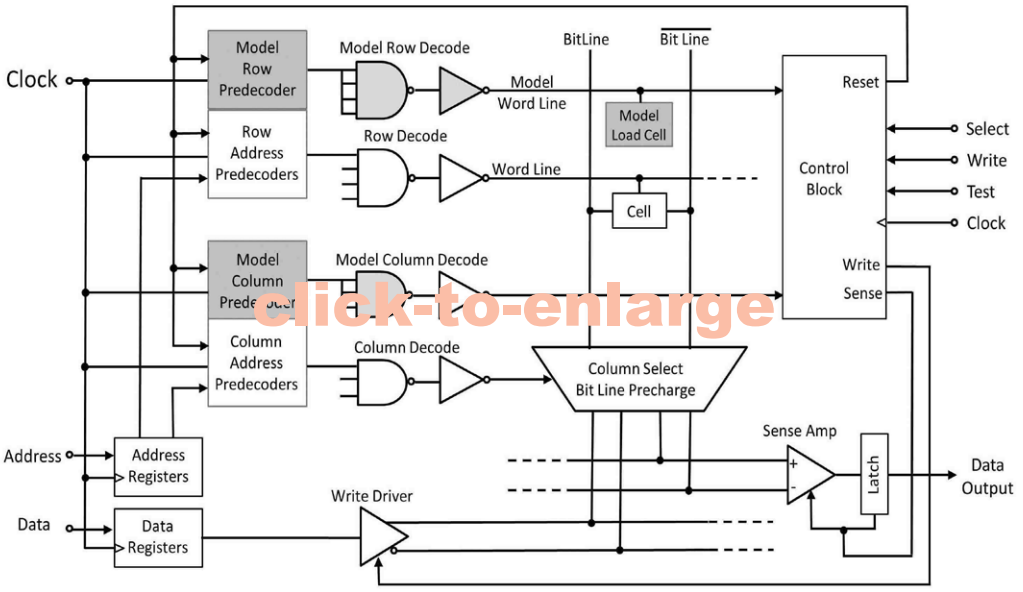

Figure 1. *Diagram of a self-timed memory.*

precoders are reset and the bit lines are precharged in preparation for the next memory access.
Memory test algorithms, like the so-called "Serial March" algorithm, are essentially a very specific sequence of writing to and reading from memory cells. For example, such a memory test algorithm may write a logic 1 into cell "a," then write a logic 0 into cell "b." If everything is OK, reading from cell "a" should result in a 1. Reading a 0 indicates a defect.
If the time that a read or write operation takes depends only on the memory itself, why do we need a high-speed clock in the first place? The speed of the clock becomes important for the MBIST logic itself. That is, it determines the speed at which the test logic can fire off these read and write commands at consecutive clock cycles towards the memory. That creates a specific sequence of operations, which forms the memory-test algorithm.
Based on this information, let's analyse the quality implications for testing a memory through MBIST running off a slow clock and then for running off a fast clock. With a slow clock, you have to look out for defects that require consecutive read/write operations fired off quickly; those defects aren't detectable by MBIST when using a slow clock. Table 1, in the online version of this article (click right) shows a list of typical defect types and the expected coverage with a slow clock. Most defects are very well covered thanks to the memory's self-timing.

Consider two examples, which are taken from the Table.
(1) A Dynamic Coupling defect is a defect where writing a value 0 to one cell forces another cell to assume a 0 as well (analogously for 1). This defect doesn't depend on how quickly the read follows the write. Thus, this type of defect is fully detectable.
(2) A Write Recovery fault occurs when a value is read from a cell just after the opposite value has been written to a cell (along the same column; and the bitline precharge has not been performed correctly). Obviously, the read operation must immediately follow the write operation. Therefore, Write Recovery faults aren't detectable using MBIST off a slow clock.

You might have noticed that the entire argument hinges off the memory's self-timing aspect. What if this self-timing logic itself is defective, such that the speed of the memory is reduced only a little, not enough to cause a catastrophic test outcome of the memory? There are three ways of answering the question. The first involves looking at the statistical defect distribution. Given the size and density of the memory array itself, the relative probability of a defect with exactly these properties in the self-timing logic is very, very small.

The second answer is that [to contrive a test that will show up this defect] you may be able to place two edges of the slow test clock so close to each other that you effectively get an at-speed clock effect lasting exactly two pulses. This will give you some confidence that the self-timing logic is operational, but the availability of this solution depends on the capabilities of the ATE (Automated Test Equipment).
Lastly, if the memory has such a speed-reducing global defect, even the simplest functional test will uncover its presence. So, we are pretty well covered at this end as well.

Special thanks to Dr. Benoit Nadeau-Dostie, Chief Architect at Mentor Graphics Corporation, for his contribution and discussion.

Martin Keim joined the Silicon Test Solutions Group of Mentor Graphics, Wilsonville, Oregon, USA, in 2001, where he is currently the Engineering Manager of the Memory Built-In Self-Test team. Dr. Keim has served for many years on the organizing committee of the International Symposium for Testing and Failure Analysis, for which he will be the General Chair in 2016.

Complete article, here

DOES YOUR OP AMP OSCILLATE?

By Barry Harvey, Linear Technology

We analogue designers take great pains to make our amplifiers stable when we design them, but there are many situations that cause them to oscillate in the real world. Various types of loads can make them sing. Improperly designed feedback networks can cause instability. Insufficient supply bypassing can offend. Finally, inputs and outputs can oscillate by themselves as one-port systems. This article will address common causes of oscillation and their remedies.

Some basics

Figure 1a shows the block diagram of a non-rail-to-rail amplifier. The inputs control the g_m block which drives the gain node and is buffered at the output. The compensation capacitor C_c is the dominant frequency response element. The return of C_c would go to ground if there were such a pin; however op amps traditionally have no ground and the capacitor current will return to one or both supplies. Figure 1b is a block diagram of the simplest amplifier with rail-to-rail output. The input g_m's output current is sent through a "current coupler" that splits the drive current between the output transistors. Frequency response is dominated by the two C_c /2s, which are effectively in parallel. These two topologies describe the vast majority of op amps that use external feedback.

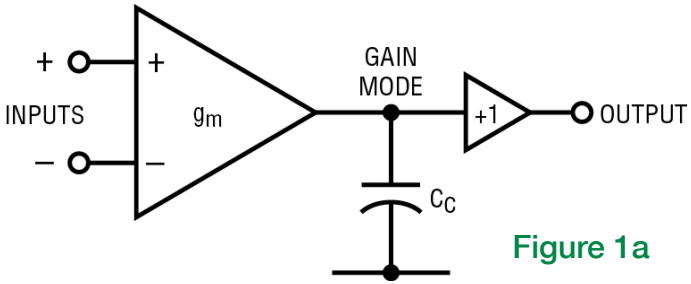

Figure 1a

Figure 1c shows the frequency responses of our ideal amplifiers, which display similar behaviour although they are electrically different. The single-pole compensation created by gm and Cc gives a unity-gain-bandwidth product frequency of GBF=g_m/(2π C_c). The phase lag of these amplifiers drops from -180 to -270° around GBF/A_{vol}, where A_{vol} is the open-loop amplifier DC gain. The phase hangs in at -270° for frequencies well above this low frequency. This is known as "dominant pole compensation," where the C_c pole dominates the response and hides various frequency limitations of the active circuitry.

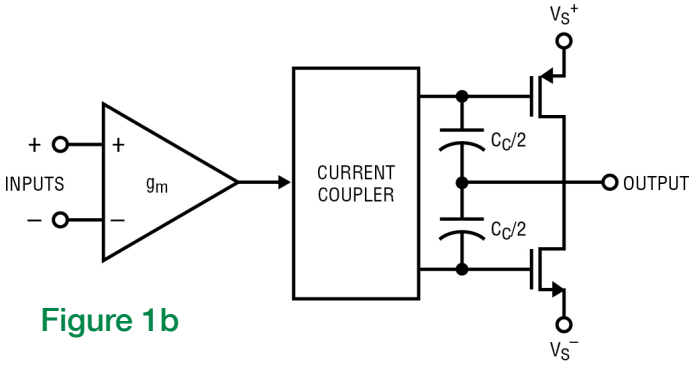

Figure 1b

Figure 2 shows the open-loop gain and phase response over frequency for the LTC6268 amplifier. This is a useful low-noise 500 MHz amplifier with rail-to-rail outputs and only 3 fA bias current, and is a good example of real amplifier behaviour. The dominant compensation's -90° phase lag starts at about 0.1 MHz, reaches -270° at about 8 MHz, but moves past -270° beyond 30 MHz. In practice, all amplifiers have high-frequency phase lags additional to the basic dominant compensation lag due to extra gain stages and the output stage. Typically, the extra phase lag starts at around GBF/10.

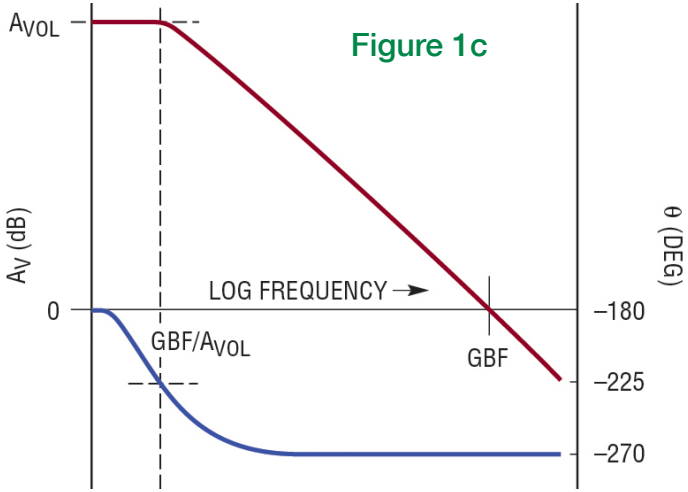

Figure 1c

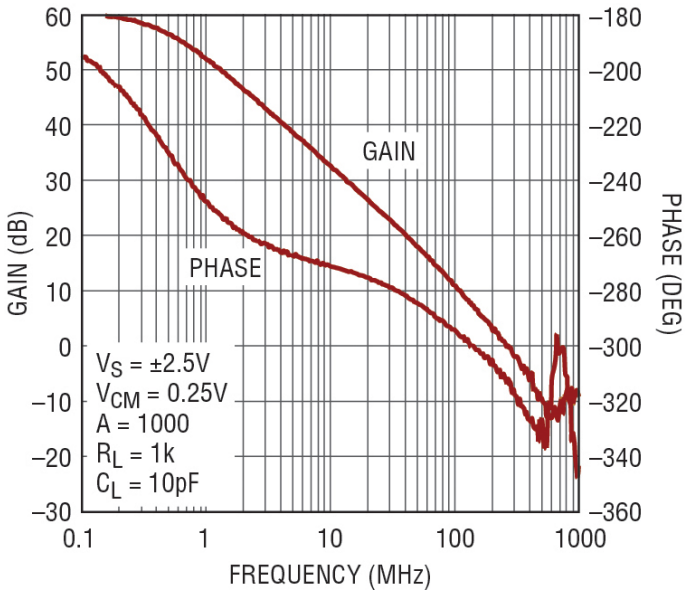

Figure 2.

Stability with feedback is a matter of 'round the loop gain and phase; or A_{vol} times the feedback factor, or loop gain in short. If we connect the LTC6268 in unity-gain configuration, then 100% of the output voltage is fed back. At very low frequencies the output is the negative of the –input, or -180° phase lag. Compensation adds another -90° lag through the amplifier, setting -270° lag from –input to output. Oscillation will occur when the loop phase lag increases to ±360°, or multiples of it, and the loop gain is at least 1V/V or 0 dB. The phase margin is a measure of how far from 360° the phase lag is when the gain is 1V/V or 0 dB. Figure 2 shows us that the phase margin is about 70° (10 pF red curve) at 130 MHz. This is a very healthy number; phase margin down to perhaps 35° is usable.

A less popular topic is gain margin, although it is just as important a parameter. When the phase descends to zero margin at some high frequency, the amplifier will oscillate if the gain is at least 1V/V or 0 dB. As shown in Figure 2, when the phase drops to 0 (or multiples of 360°, or -180 as in the figure), the gain is about -24 dB at approximately 1 GHz. This is a very low gain; no oscillations will occur at this frequency. In practice one wants at least 4 dB of gain margin.

This very comprehensive article continues to discuss; Decompensated amplifiers; Feedback networks; Load issues; Strange impedances; and the effects of Power supplies. The author concludes that the designer needs to consider parasitic capacitance and inductance associated with each op amp terminal, and the nature of the load. The amplifiers are designed to be stable within a nominal environment, but each application requires its own analysis. Download the complete article pdf.

THE FULLY DIGITAL RADIO TRANSMITTER: IS IT REAL?

By Steve Taranovich, EDN

In this article, Steve Taranovich takes an overview of several recent attempts to create an "all-digital" radio, prompted by a recent announcement from Cambridge Consultants of an FPGA-based architecture.

Cambridge Consultants are claiming the world's first fully digital radio transmitter built only from computing power. There are no analogue components such as a high-speed D to A converter with amplifier, although I would think they would need a Power Amplifier (PA) to broadcast a great distance. This is a Digital Radio transmitter and not part of a Software Defined Radio (SDR) architecture which requires analogue components. Cambridge Consultants demonstrated the transmitter at Mobile World Congress (MWC) in Barcelona, in early March.

The radio has been named Pizzicato: my Italian wife tells me that Pizzicato means pinched or to pluck as in a stringed instrument. More about this new architecture later; now let's see what came before this effort leading up to the Cambridge solution.

Many so-called All-Digital Radios have been tried in the past. Here are some that stand out, but in my sceptical analogue brain I find it hard to conceive a truly All-Digital Radio, only possibly some main architectural sections as is evidenced in the following examples from some EEs in IEEE tech papers.

The pulsed UWB transmitter

In Reference 1 we see a 2008 design for pulsed-ultra-wideband transmitter architectures. The developers have two designs: delay line-based and ring oscillator-based. To illustrate the advantage of a digital format architecture, the author shows that the two architectures are synthesised and place-and-routed

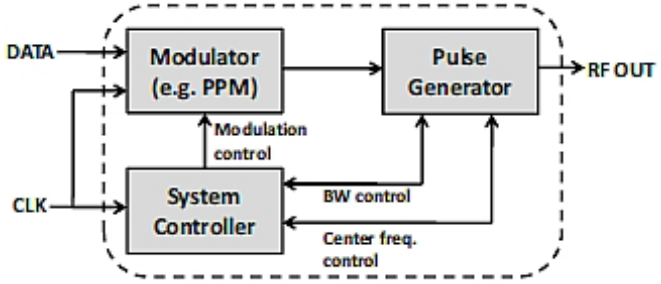

Figure 1. *Reference 1 architecture for the "All-digital" pulsed UWB transmitter. (Image courtesy of Reference 1)*

(PAR) using existing design tools - no custom routing or care as is typical in analogue layout and routing. Another advantage of a digital system is that when the desired centre frequency and bandwidth are slightly off due to variations induced by PAR, those frequency-related parameters can be easily tuned digitally. See Figure 1 for the architecture block diagram.

Not bad, but not my idea of a fully-flexible all-digital radio design.

Radio transmitter architecture with all-digital modulator

Reference 2 is a mid-2000 design with somewhat limited application to an "Opportunistic Radio" which shares spectrum in licensed and unlicensed bands for secondary users. The architecture used to address the need for a flexible and reconfigurable radio type is a direct modulation architecture based on an all-digital I&Q modulator followed by a linear power amplifier. With a digital modulator, the D to A conversion is able to take place at the speed of RF, so at no less than 2X the channel frequency of the signal in order to meet the Nyquist criteria. See Figure 2.

The "all-digital" modulator is based upon the sigma-delta filtering plus a set of analogue band-pass BAW-CRF (bulk acoustic wave-coupled resonator filter) filters where only one

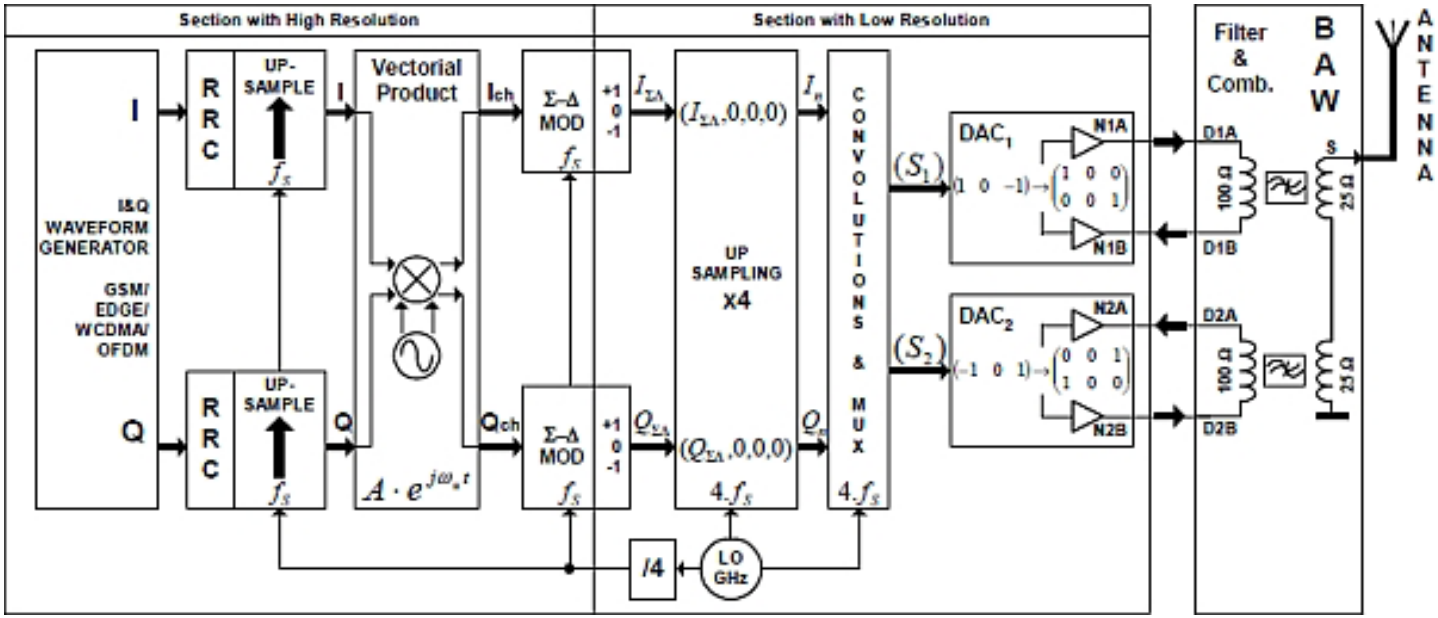

Figure 2. *The "All-Digital" Radio Transmitter Architecture based on sigma-delta modulation. (Image courtesy of Reference 2)*

path to the filter is active at any one time. The D to A converter is clocked by an "all-digital" RF synthesizer. There is no Power Amplifier (PA) used so this solution is very linear and good for a Zigbee protocol with low power and shorter range. Noise shaping is done at higher frequencies similar to the original sigma-delta designs that handles lower audio-type frequencies.

In most of these examples, we look five or ten years in the future to the maturity of such architectures. The maximum speed of the sigma-delta modulators needs to be increased, higher power levels and the PA issue needs to be addressed, and power system efficiency needs to be improved.

FPGA-Based All-Digital Transmitter with RF Output for Software Defined Radio

This mid-2000s effort used an FPGA that directly synthesises the RF signal in the digital domain which eliminates most analogue and RF components. See Figure 3.

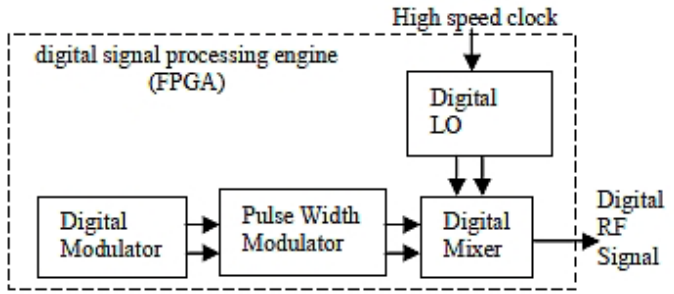

Figure 3. *The All-Digital transmitter architecture. (Image courtesy of Reference 3)*

This is another architecture that uses digital architecture as well as a band-pass delta-sigma modulation scheme to generate binary signals at radio frequency. Binary signalling coupled with a switch mode PA will result in higher efficiency than conventional PAs. One shortcoming of this architecture is that the band-pass delta-sigma modulator has to operate at four times the output centre frequency, placing it in the multi-GHz region in most cases.

In the continuation of this article, Steve notes several more alternative approaches to an all-digital radio, and concludes with some further observations on the Pizzicato report. Click to download the pdf.

LOW-POWER, LOW-COST, DIFFERENTIAL AMPLIFIER WITH HIGH CMRR

BY CHAU TRAN AND JORDYN ROMBOLA, ANALOG DEVICES

Many applications require special differential amplifiers that achieve both high performance and low power dissipation. A simple solution is to use a dual non-inverting precision amplifier with a resistive gain network, but the DC common-mode level of the differential outputs will depend upon the common-mode voltage at the input. Figure 1 shows the ADA4805-2 high-performance dual amplifier configured as an inverting differential amplifier. This low-noise circuit operates with common-mode gain of +1 from the two non-inverting inputs to the outputs, enabling easy common-mode control for single-supply operation. The low quiescent current (500 µA per amplifier) suits the circuit for low-power, high-resolution data-conversion systems.

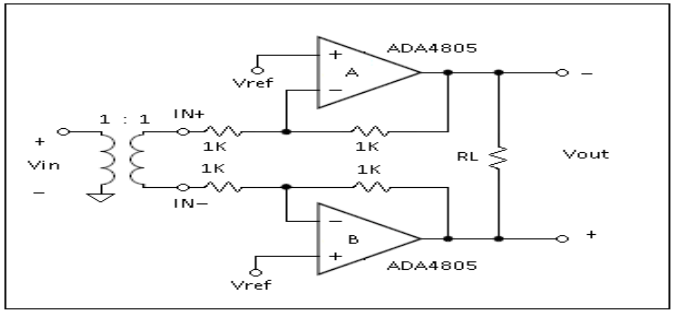

Figure 1. *Inverting differential amplifier configuration.*

This circuit offers a combination of low distortion and low quiescent current. The dual op amp solution results in lower system cost, while the differential amplifier topology yields better performance.

High-resolution ADCs often have differential inputs to reduce common-mode noise and harmonic distortion. Therefore, single-ended input voltages must be converted to differential input voltages before they can be useful for ADCs. This circuit can be used to convert a single-ended signal into a differential signal to drive the ADCs while the common-mode voltage can be set by the ADC's reference.

The transformer converts the single-ended signal into a differential signal. With a system gain of –1, the differential output can drive the ADC. The common-mode output voltage is determined by the voltage Vref at two non-inverting inputs.

To eliminate the DC common-mode at the inputs, add a series capacitor to each input, or a transformer as shown in Figure 1. Also, the differential input signals can be applied directly to IN+ and IN– pins. In this case, the gain resistors become part of the input resistance for the source.

The differential-mode frequency response plot was taken by applying a differential input voltage at the two input terminals (IN+ and IN–). The gain resistors become part of the input resistance for the source. This provides better noise performance than the non-inverting configuration but limits the flexibility in setting the input impedance separately from the gain.

For the common-mode frequency response plot, the common-mode voltage is applied at two non-inverting input terminals (Vref) with IN+ and IN– tied together.

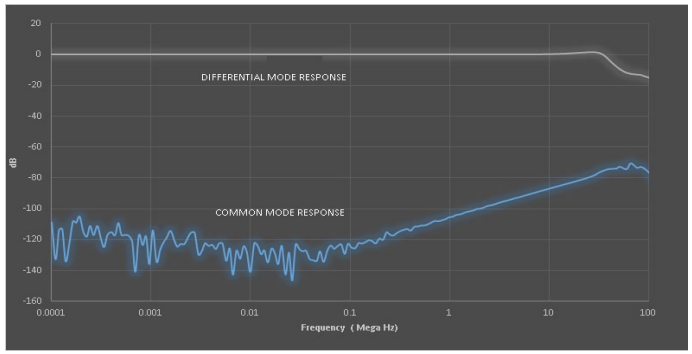

Figure 2. *Performance of the inverting differential amplifier.*

Chau Tran *[chau.tran@analog.com] joined Analog Devices in 1984 and works in the Instrumentation Precision Technology (IPT) Group in Wilmington, MA.*

Jordyn Rombola *[jordyn.rombola@analog.com] is a product engineer in the Instrumentation Precision Technology (IPT) Group.*

designideas

- Transient load gives power systems a workout
- Monitor a differential pair for valid & fault conditions
- NRZ to AMI converter uses single supply

Transient load gives power systems a workout

By Mike Rose

Using the fast dynamic load described in this Design Idea to test the transient response of a power system can reveal many critical operating characteristics. The voltage deviation resulting from a fast current step can provide insight into a regulator's phase margin (Reference 2). Also, for power supplies that are located some distance from the point of load, transient testing can help determine the effective series interconnect inductance, shunt capacitance, and ESR. While the phase margin for commercial power supplies is typically verified by the supplier, adding remote sense can often destabilise the supply. The interconnect inductance and the load capacitance can introduce additional phase shift into the regulator control loop feedback, compromising stability. Many engineers have seen the tell-tale low frequency sine wave riding on the regulator output.

Performing transient testing on an assembled system allows a quick check on the system's dynamic regulation stability and accuracy (Figure 2). Most commercial dynamic electronic loads have fairly slow current slew rates, which tend to limit their usefulness in testing faster regulator control loops, which can often reach steady state within 50 µsec or less following a large load transient. A current slew rate of 10 A/µsec or higher is needed for many high power systems.

Figure 1 is an adaptation of an application note (Reference 1) with a few notable improvements. The maximum power level has been increased to 150W, and it has been designed specifically for 3.3V, 5V, and 12V regulator outputs. R1-R3 form a resistive load switched by a single N-channel low-side MOSFET. The load resistors can be sized and populated to realise a large number of possible load combinations.

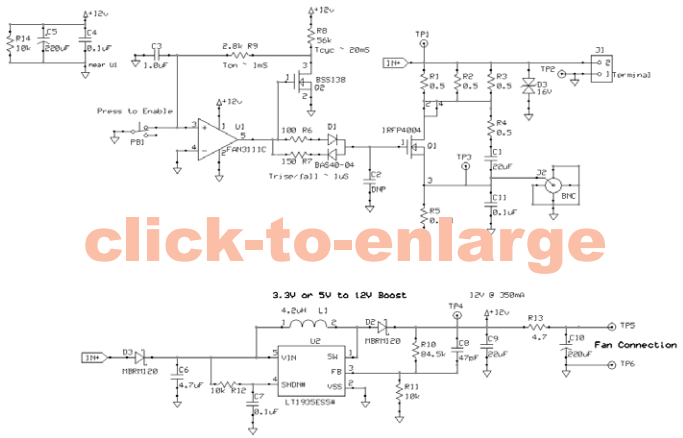

At the core of the circuit, MOSFET driver U1 with a Schmitt trigger input is used to drive the MOSFET, and forms a relaxation oscillator with Q2, R8, R9, and C3. With the component values shown, the duty cycle is approximately 5% at a cycle time of 20 msec. Having a relatively low duty cycle permits a modest cooling solution.

R6 and R7 independently adjust the rise and fall times in combination with the MOSFET's input capacitance, C_{ISS}. With the values shown, rise and fall time are approximately 1 µsec. At this slew rate, the peak MOSFET gate current is around +110/-75 mA, which is well below the maximum rating of 1.4A for U1. C2 can be added to further slow the edge rate. With a 1 µsec rise/fall time and the relatively large gate resistors adding dissipative damping, MOSFET gate switching resonances will not be significant. R4 and C1 help dampen line resonances when the MOSFET switches off. The value of R4 is determined from the effective line inductance and the input capacitance. The value of 0.5Ω has proved to be effective for typical wiring scenarios.

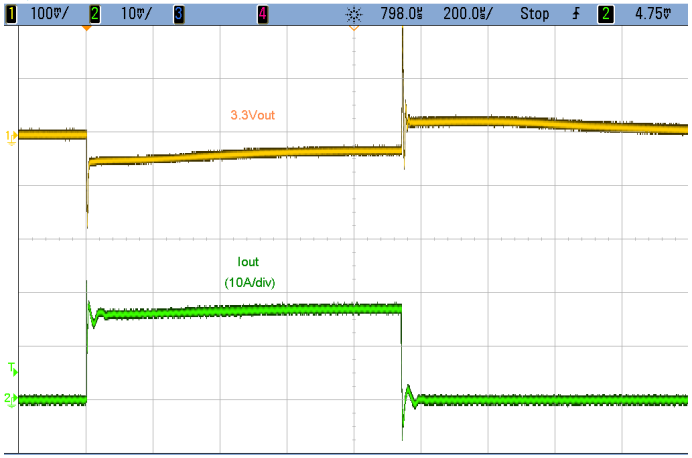

Figure 2. *Application of a 50W load transient to a commercial power supply exposes the response of the supply and the interconnect characteristics. From the response, approximate phase margin can be determined (Reference 2), along with regulator loop bandwidth and accuracy.*

One of the more convenient features of this implementation is the two-wire connection to the DUT. For 3.3V and 5V systems, a 12V boost converter is included to power the MOSFET driver and gate. No other connections or power sources are needed. The output of the boost converter can source approximately 350 mA with a 3.3V input which may limit the amount of current that can be provided to charge the MOSFET's gate. A low ESR aluminium capacitor, C5, supplies some of the initial gate charging current for faster current edge rates. For 12V operation, a direct connect version can be assembled by replacing C7 with a 0Ω resistor, disabling the boost converter. There will be some voltage drop across L1 and D2, but this will not affect proper operation of the circuit. D3 protects the boost converter circuits from reverse polarity (note that the body diode of Q1 will conduct during reverse polarity, likely leading to excessive dissipation in the MOSFET).

The entire circuit fits comfortably on a 3 × 5 in. two-layer PCB, including the heat sink and a small 12V fan. With only two wires to connect, operation is very straightforward. The tester's leads must be short and have low inductance to prevent ringing from the lead reactance. DUT connection should be made near the point of load or the remote sense location. The tester's common and the voltage probe return leads should be connected at a single location. This location should also be selected to have a low impedance path back to the power source.

Pressing the momentary pushbutton PB1 starts the astable circuit and the dynamic load begins switching. If desired, quiescent PSU loading can be provided externally. R5 and J2 provide a convenient high bandwidth means for measuring the pulsing current. A straight 50Ω coax can be connected directly to an oscilloscope input for monitoring current with a scale of 1 mV/A.

The voltage measurement should also be made near the point of load or remote sense point and should be AC-coupled to a second scope input. The voltage probing must be done with care. Probe inductance from a distant ground/return lead will cause misleading measurements. A small series resistor (several ohms) at the probe tip can be added to dampen out the high frequency ringing from the probe's ESL. Also, avoid probing right at the pads of a very low ESR decoupling capacitor which may filter or dampen the voltage response unrealistically.

REFERENCES

Reference 1. Application note 1716.0, Intersil, Paul Traynham and Dan Swank, January 26, 2012

Reference 2. Application note AN1733, Texas Instruments, SNOA507, November 2007

Mike Rose is a consulting engineer specialising in power, analogue, signal integrity and embedded computing.

Monitor a differential pair for valid & fault conditions

By Budge Ing

It is advantageous to know if there is a valid signal on a differential signal bus. This Design Idea detects differential data transmission and reports a loss of signal (LOS) to a microcontroller or other monitoring device.

The circuit of Figure 1 detects a continuous differential signal from -7V to +12V with a minimum amplitude difference of 200 mV. The circuit detects open inputs, shorted inputs, and inputs stuck at high or low. The circuit's output is low when any one of these conditions occurs, and high when there is a differential pulse train. Only a single 5V supply is needed.

U1 (MAX3280), an RS-485 receiver, accepts differential signals of 200 mV minimum at A and B, where the voltage can be from -7V to +12V. Pin RO of U1 is high in the absence of input signals, because U1 is a true fail-safe device that provides a high output when the inputs are open or shorted. RO can also be high or low when one of the inputs is stuck high and the other is stuck low. R1 and C1 form a lowpass filter to average U1's output into a DC voltage when the output at RO is a pulse train.

U2 (MAX992), a dual open-drain comparator, is configured as a window comparator, with R2-R4 setting the voltage thresholds to 0.25V and 4.75V. Only when the filtered voltage of a pulse train falls between the two thresholds will the final OUT signal be high.

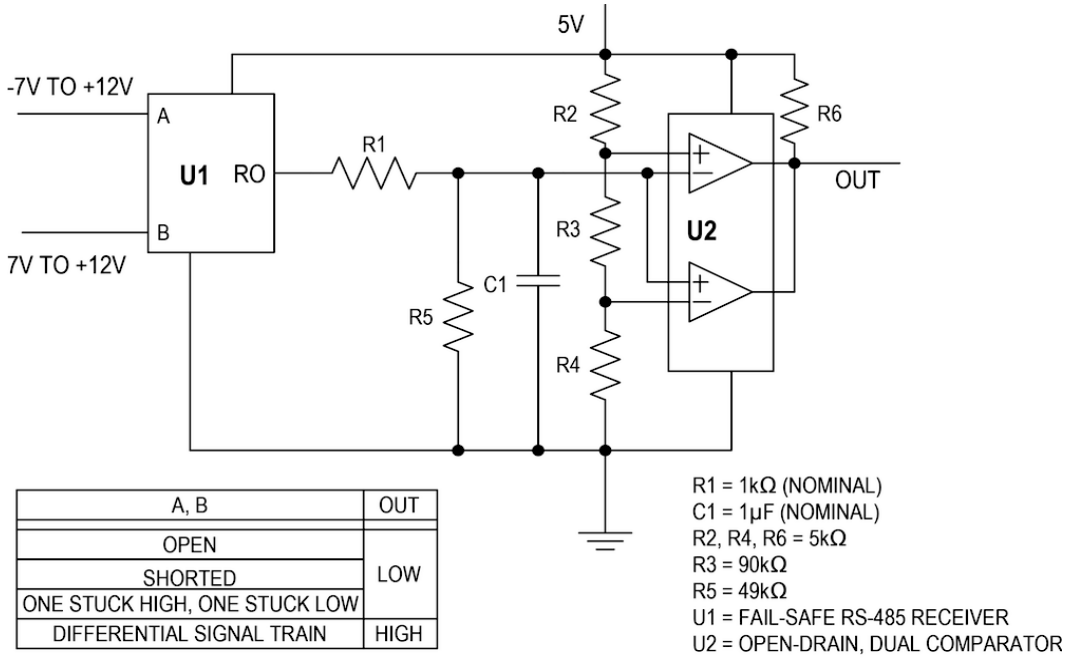

A, B	OUT
OPEN	LOW
SHORTED	
ONE STUCK HIGH, ONE STUCK LOW	
DIFFERENTIAL SIGNAL TRAIN	HIGH

Figure 1. *Circuit detects a continuous differential signal*

Budge Ing is an Applications Engineer at Maxim Integrated.

NRZ to AMI converter uses single supply

By Ramachandra Mutagi

Alternate mark inversion (AMI) coding is often used in digital data transmission over cables as it has no DC component. Besides, the bandwidth of an AMI signal is lower than the equivalent return-to-zero (RZ) format. Normally, to generate a bipolar waveform such as AMI, one needs positive and negative supplies. Also, the circuit will likely employ analogue components for bipolar waveform generation. However, this Design Idea eliminates all these requirements and generates an AMI waveform from an NRZ input using only a few gates, a flip-flop, and a single 5V supply.

Referring to Figure 1, the NRZ signal (Figure 2a) is gated with the clock using the AND1 gate, generating an RZ waveform (Figure 2b). The RZ signal clocks a D-FF connected as a frequency divider. The RZ signal is then gated with the Q and /Q outputs of the flip-flop, separating alternating pulses on two lines at the outputs of the AND2 and NAND gates. The NAND gate is used on the second line to obtain an inverted waveform (Figure 2c).

Since the delay of the NAND is more than the AND gate, AND4 is used at the output of AND3 to compensate (this can be changed depending on the logic family used). The outputs of the AND4 and NAND gates drive 75Ω resistors which effectively add the voltages at the gate outputs. If both the outputs are high, the voltage at the junction of the resistors is a high level. If one of the outputs is low and other high, the voltage at the junction is half the high level voltage. When both outputs are low, the voltage at the junction is close to 0V. Thus, the waveform at the junction of R1 and R2 shows positive and negative pulses around a DC level. This signal passes through DC block C1, and we get a true bipolar waveform with zero DC level at the output (Figure 2d).

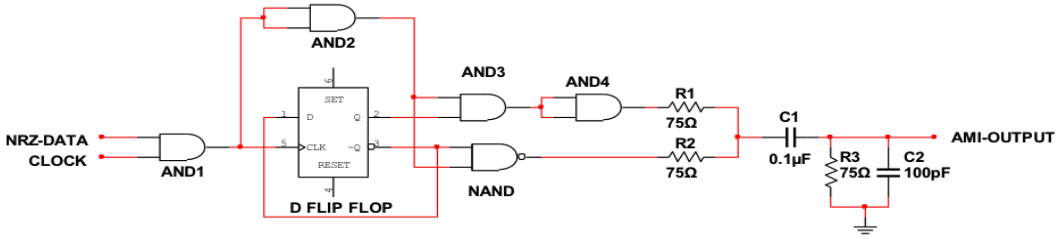

Figure 1. *The NRZ to AMI converter generating bipolar pulses uses single supply*

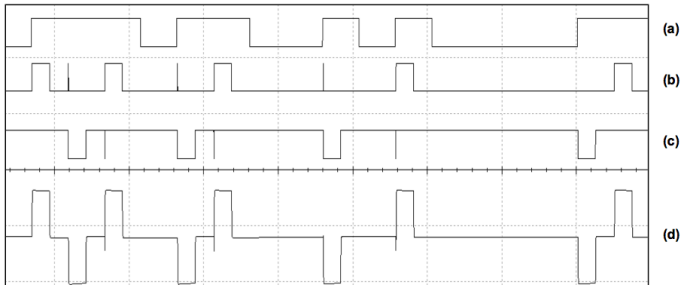

Figure 2. *The waveforms: (a) NRZ input; (b) clock gated RZ output at AND1; (c) NAND output; (d) AMI output*

Figure 2 shows waveforms of the simulated circuit. The simulator does capture very small spikes appearing at the gate outputs, which will not cause problems in real use. The NRZ signal was generated at 2.048 Mb/sec. Because of the TTL devices and 5V supply, the peak-peak signal level is less than ±2.5V. If higher amplitude is needed, CMOS devices with higher logic swings may be used.

R. N. Mutagi is a Professor and Head of Electronics and Communication Engineering Department at Indus University, Ahmedabad, India. Prior to joining the academia in 2010 he served with the Aerospace companies in Canada and Indian Space Research Organization (ISRO) as R&D Engineer and Manager.


productroundup
SL871
Telit
TOSHIBA
TPHR8504PL
MOSFET
5.0mm
IOR
07BFA6
EXAR
XRP9710
XRP9711
efus
Linux
TOSHIBA
TB67H303HG
econais
EC19D
lesswire
WiBear 11ac
ROHM
SEMICONDUCTOR
LINEAR
TECHNOLOGY
LT8490
MICROCHIP
dsPIC33EP128GM304
DIODES
AP7340
IOR

productroundup

Harvested-energy storage capacitors, in distribution

Distributor Farnell element14 now has Vishay's hybrid storage 196 HVC ENYCAP capacitors which are portable, compact and facilitate high-energy storage. The capacitor range has the highest energy density (13 Wsec/g) as well as including polarised storage capacitors with high capacitance (up to 90F). Suitable as a storage device for energy harvesting, the capacitors are an alternative to rechargeable backup batteries as they can hold 70% of their charge for as long as a month. They support voltages ranging from 1.4V (single cell) to 2.8V / 4.2V / 5.6V / 7.0V / 8.4V (multiple cells) and are available as stacked through-hole, surface mount flat and lay-flat configurations.

Complete article, here

60V, 3x output, 29 µA Iq buck/buck/boost DC/DC

Designed to maintain regulation in automotive and heavy truck start/stop systems, the LTC3899 is a triple output (buck, buck, boost) 29 µA quiescent current, synchronous DC/DC controller that maintains output voltage regulation over a 2.2V to 60V input range. A 12V car battery can droop to less than 4V during engine restart or cold crank, causing reset of infotainment systems and other electronics that operate from 5V and higher. The boost converter feeds two step-down converters, avoiding output voltage dropout when a car battery droops, a useful feature in automotive start/stop systems that shut off the engine at idle to save fuel.

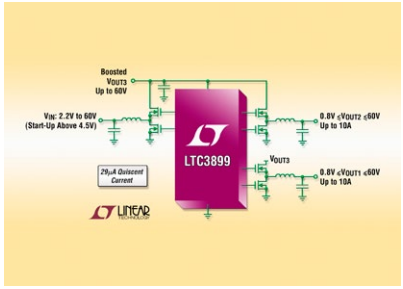

Complete article, here

MEMS sensors have integrated MCUs for Android phones

BHI160 and BHA250 integrate Bosch Sensortec's lowest power sensor hub and sensors, using less power for always-on sensor applications such as fitness tracking, indoor navigation and gesture recognition. The devices implement the full Android Lollipop sensor stack with lowest power sensor and sensor hub combination, reducing power consumption by up to 95% for always-on applications in mobile devices by handling sensor processing and data batching. BHI160 and BHA250 offload sensor processing from an application processor, as well as buffering sensor data locally on the devices; devices integrate a 3- or 6-axis MEMS sensor with the Bosch Sensortec DSP "Fuser Core".

Complete article, here

Development platform supports FTDI's microcontroller architecture

MikroElektronika, in partnership with FTDI Chip, has produced a complete development package to support the FT90X series of 32-bit application-optimised microcontroller units (MCUs). FTDI Chip developed the FT90X 32-bit architecture as an efficient and minimum-footprint processor architecture that is focussed on on the connectivity and HMI-related tasks that the company specialises in. The MCUs in the series have ample headroom, FTDI executives have commented, to run users' application code in addition, or to be used as a stand-alone MCU platform. mikroC, mikroBasic and mikroPascal compilers fully support the FT90X.

Complete article, here

productroundup

Sponsored Content

MSP430I2040

Design applications like factory and building automation, energy metering and industrial lighting using the new MSP430i204x industrial microcontrollers. MSP430i204x microcontrollers are integrated with smart analog for high accuracy, precision and cost savings. These devices also meet the broad temperature range requirements from -40 to 105 °C needed for industrial and smart grid applications. They consist of a powerful 16-bit RISC CPU, a DCO-based clock system that generates system clocks, a power management module (PMM) with built-in voltage reference and voltage monitor, two to four 24-bit sigma-delta ADC converters, a 16-bit hardware multiplier, two 16-bit timers and up to 16 I/O pins.

Complete article, here

MSP430FR6989

The MSP430 ultra-low-power (ULP) FRAM platform combines uniquely embedded FRAM and a holistic ultra-low-power system architecture, allowing innovators to increase performance at lowered energy budgets. FRAM technology combines the speed, flexibility, and endurance of SRAM with the stability and reliability of flash at much lower power.

The MSP430 ULP FRAM portfolio consists of a diverse set of devices that feature FRAM, the ULP 16-bit MSP430 CPU, and intelligent peripherals targeted for various applications. The ULP architecture showcases seven low-power modes, which are optimized to achieve extended battery life in energy-challenged applications.

Complete article, here

InstaSPIN enabled C2000™ Piccolo LaunchPad

The InstaSPIN™-MOTION (and InstaSPIN-FOC) enabled C2000™ Piccolo LaunchPad is an inexpensive evaluation platform designed to help you leap right into the world of motor control using the InstaSPIN-MOTION or InstaSPIN-FOC solution. The LaunchPad is based on the Piccolo TMS320F28069M with unique features such as 256KB of on board flash, 12bit ADC, I2C, SPI, UART, CAN, dual Encoder support and InstaSPIN libraries in on-chip execute only ROM memory. The LaunchPad includes many board hardware features such as an integrated isolated XDS100v2 JTAG emulator for easy programming and debugging, Works with various BoosterPacks, but specifically for BOOSTXL-DRV8301.

Complete article, here

Delfino TMS320F2837xD

The Delfino TMS320F2837xD is a powerful 32-bit floating-point microcontroller unit (MCU) that lets designers consolidate control architectures and eliminate multiprocessor use in high-end systems. The dual real-time control subsystems are based on TI's 32-bit C28x floating-point CPUs, which provide 200 MHz of signal processing performance in each core. The C28x CPUs are further boosted by the new TMU accelerator, which enables fast execution of algorithms with trigonometric operations common in transforms and torque loop calculations; and the VCU accelerator, which reduces the latency for complex math operations common in encoded applications.

Complete article, here

productroundup

Analysis tool stops software "erosion"; supports MISRA C

Enabling developers to write better embedded software is the goal of the creators of the Bauhaus Suite, a software development environment from tool vendor Axivion GmbH. The tool suite supports MISRA C:2012 and allows automated static analysis for C# programs integrated into Microsoft Visual Studio. The Axivion Bauhaus tool suite, available now in release 6.2, targets safety-relevant software in automotive and medical environments. Supporting the MISRA C:2012 standard, the tool tests embedded software and provides information to the programmer as to in which lines its code deviates from the rules as set by the relevant standards. It also provides hints to the programmer which program statements need to be corrected or at least described in a comment. The selection of programming languages supported has been enhanced with a view to GUI design in industrial environments. As a result, the tool enables users to perform in-depth static analysis of C# code.

Complete article, here

USB vector analyser: portable passive characterisation

Anritsu's ShockLine single-port USB vector network analysers offer an economical and compact route to testing cables, antennas, and other passive RF devices. The MS46121A series of 1-port USB VNAs embed the company's ShockLine VNA technology in a compact package. Two models with frequency coverage of 40 MHz to 4 GHz and 150 kHz to 6 GHz are powered and controlled via a user's computer to simply and cost-efficiently test passive RF devices where 1-port measurements up to 6 GHz are required. Sweeps can be conducted at 100 microseconds per point and the MS46121A VNAs have a corrected directivity of 42 dB.

Complete article, here

Differential current transformer for current measurements up to 10A

This differential current transformer is capable of carrying out current measurements at up to 10A, and is part of Harting's family of current sensors for industrial applications. The differential current transformer permits the measurement of the current deviation in supply and return conductors or the monitoring of earth leakage currents in single or multi-phase power cables. The transformer is designed for a maximum cable diameter of 110 mm, and is suitable for the detection of primary nominal currents of up to 10A with a conversion ratio of 1:600. It can be used for measurements over the frequency range from 5 to 400 Hz.

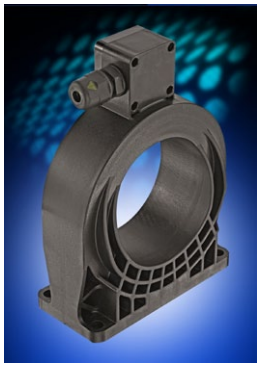

Complete article, here

Smartwatch reference model has motion and activity sensors

Toshiba Electronics Europe (TEE) has a Smartwatch Reference Model, with enhanced sport, activity and lifestyle features; the TZ1001 Smartwatch Reference Model uses Bluetooth Low Energy connectivity and Qi Wireless Charging. Built around the TZ1001MBG ApP Lite application processor, it is fitted with an accelerometer, gyroscope, magnetometer and optical pulse sensor, to measure motion, movement and heartbeat. The Reference Model provides a structure for prototyping devices aimed at the active-lifestyle and wellbeing sectors.

Complete article, here

productroundup

3W class-D audio amp drives speakers or class-AB headphones

PAM8009 from Diodes Incorporated is a 3W stereo class-D audio amplifier and class-AB headphone driver. Featuring power limit technology for speaker protection and a DC volume control for precise 64-step gain adjustment, this amplifier minimises external components to achieve a highly integrated solution in 20-pin QFN or 24-pin SOP packages. With an output power rating of up to 2.4 W/channel into 4Ω bridge-tied speakers at a THD+N of 1% (up to 3W at 10% THD+N), the PAM8009 provides a high-power, low-cost capability for driving typical small internal speakers.

Complete article, here

Sensirion reveals multi-pixel gas sensor

Sensirion announced it had developed a multi-gas sensor that it was planning to sell to smartphone makers back in November 2013. The company also offers smartphone makers a barometric pressure sensor suitable for use for indoor navigation applications. Sensirion (Staefa, Switzerland) has focused on humidity and temperature sensors but is now expanding its range of environmental sensors to include gas and pressure sensors. The gas sensor is the first in the world to be based on multiple gas-detecting pixels. This allows the sensor to perceive its surroundings using various receptors that using pattern recognition is able to interpret different concentrations on different sensors in terms of type and concentration of a wide range of gases. The sensor is housed in a package measuring 2.45 by 2.45 by 0.75 mm.

Complete article, here

MCUs with security features protect app. code and data

NXP's LPC18Sxx and LPC43Sxx microcontroller families add hardware-accelerated encryption for secure boot and secure messaging, in connected applications. Both LPC18Sxx and LPC43Sxx families integrate an AES-128 encryption engine; two 128-bit non-volatile OTP memories for encrypted, hardware-randomised key storage to prevent cloning; a true random number generator for unique key creation; and boot ROM drivers supporting secure boot. The MCUs in both families use ARM Cortex-M cores (Cortex-M3/LPC18Sxx, LPC-43Sxx/Cortex-M4 & M0) ensuring bandwidth for fast bulk data encryption or decryption without slowing down communications.

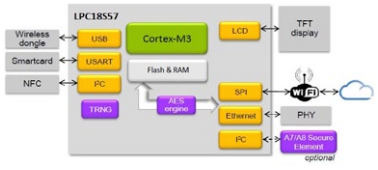

Complete article, here

Micropower wake-up timer for power-sensitive applications

Featuring pushbutton control, LTC2956 is a wake-up timer that manages 1.5V to 36V system power, that is presented as offering an improved solution compared to using microcontrollers or real-time clocks to track timing, which, Linear Technology says, overcomplicates designs. The LTC2956 periodically "wakes up" and turns on a system to perform these tasks, then turns the system off to conserve power. While "sleeping" the LTC2956 needs only 800 nA of quiescent current from a battery or rail. The wake-up period is resistor-adjustable from 250 msec to 39 days, and requires no code.

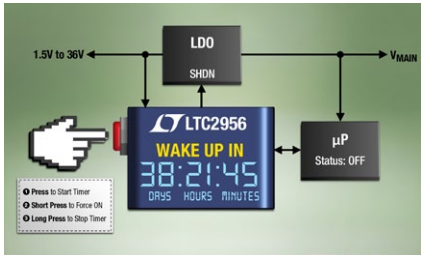

Complete article, here

productroundup

Renesas' RX700 MCU offers 240-MHz, run-from-flash operation

The RX700 microcontroller series expands the 32-Bit RX family with a part that has 240 MHz operation, security features and 4 MB of memory: a combination that Renesas presents as an increase in productivity and security for industrial connected-device applications. Developed for use in industrial equipment, the new series doubles the CPU operating frequency to 240 MHz from the 120 MHz of previous products and is

available with up to 4 MB of on-chip flash memory. The flash memory is built in Renesas' own MONOS, 40-nm process which achieves 8-nsec access times, four times faster than alternative technologies, and with long endurance and reliability.

Complete article, here

ARM-core microcontroller for secure, connected, high efficiency motor control

Providing secure connectivity at all levels, Freescale's Kinetis KV5x MCU family uses the full performance potential of the ARM Cortex-M7 core in the rapidly evolving digital motor control market. The migration toward digital-based control systems with secure networking capabilities is addressed by the Kinetis KV5x MCU which incorporates an IEEE 1588 Ethernet controller, a cryptographic acceleration unit with random number generator, and a memory protection unit. With motors often employed in

safety-critical environments such as manufacturing process control, these features allow developers to implement new services via the IoT infrastructure while protecting against erroneous inputs that could lead to an undesired operating condition.

Complete article, here

Runtime visualisation tool adds Segger embOS and J-Link support

Percepio AB, the Swedish developer of RTOS visualisation tools, has added embOS-Trace. This tool offers a new level of insight into the runtime world for embedded software developers using Segger embOS and Segger J-Link debug probes. Percepio's visualisation offers over 20 views, which are interconnected to allow viewing the trace from multiple perspectives, for instance CPU usage, memory allocation, task synchronisation and communication events, as well as selected state variables and input/output values. The

insight facilitates debugging, validation, profiling, documentation and training avoiding weeks of troubleshooting and increasing development speed.

Complete article, here

Bluetooth SoCs at "ultra-low" power in Silicon Labs' Blue Gecko identity

Silicon Labs has introduced a Bluetooth Smart product portfolio designed to help developers minimise the energy consumption, cost and complexity of wireless IoT designs. It comprises Blue Gecko Modules, SoCs, Development Kit and Software Stack. Blue Gecko SoCs are based on the ARM Cortex-M3 and M4 cores and offer 128 to 256 kB flash sizes and 16 to 32 kB RAM sizes. The SoCs integrate an array

of low-energy peripherals as well as Silicon Labs' Peripheral Reflex System (PRS) for autonomous peripheral operation. The Blue Gecko SoC family also offers a roadmap of enhanced flash and RAM memory sizes.

Complete article, here

7 TIPS FOR CREATING A RELIABLE EMBEDDED SYSTEM

BY JACOB BENINGO

Despite the hopes and dreams of many embedded engineers, reliable code doesn't happen by accident. It is a painstaking process that requires developers to maintain and manage every bit and byte of the system. There is usually a sigh of relief when an application is validated "successfully" but just because the software is running correctly in that moment under controlled conditions doesn't mean that it will tomorrow or a year from now. There is a plethora of techniques for creating a reliable embedded system, ranging from a well-disciplined development cycle through strict implementation and system checking. An entire library could easily be filled with books on reliable software design. But there are seven tips that are easily implemented that can go a long way to ensure that a system performs more reliably and catches unexpected behaviour.

Tip #1 - Fill ROM with known value

Software developers tend to be a very optimistic group, at least as far as their expectations of how faithfully their microcontroller will run their code over time. The thought of the microcontroller jumping out of the application space and executing in unintended code space seems like a fairly rare case; however, the opportunity for this to occur is no more than a buffer overflow or the dereferencing of a faulty pointer away. It can and DOES happen! The resulting behaviour of the system would be undefined since memory could have all 0xFF's in the space by default or, since the region of memory normally isn't written, the values could have decayed into almost anything.

There is a pretty neat linker or IDE trick, though, that can be used to help identify and recover the system from just such an event. The trick is to use the FILL command to fill unused ROM with a known bit pattern. There are many different possible combinations of what can be used to fill the unused memory with but if the intent is to build a more reliable system the obvious choice is to place an ISR fault handler in this location. If something goes wrong and the processor starts to execute code outside of program space then the ISR will fire, providing the opportunity to store the state of the processor, registers and system before deciding on a corrective course of action.

Additional information on how to use FILL and alternative strategies for its use can be found in "Improving Code Integrity Using FILL" located here.

Tip #2 – Check Application CRC

One of the great benefits available to embedded engineers is that our IDEs and tool chains can automatically generate an application or memory space checksum from which the application can be verified. The interesting thing is that in many of these cases the checksum is used only at the time of loading program code onto a device.

If a CRC or checksum is kept in memory, though, then verifying that the application is still intact at start-up (or even periodically for long running systems) is a great way to ensure that something unexpected won't occur. Now the chances that a programmed application will change is small, but considering the billions of microcontrollers shipped each year and the possible harsh operating environments, the chances of a corrupted application is not zero. Even more likely is that a bug in the system could cause a flash write or flash erase in a sector, resulting in a corrupted application.

Read tips 3 to 7 in the complete article; download the PDF.

EDN europe

EDN-EUROPE is published 11 times in 2015 by
European Business Press SA,
7 Avenue Reine Astrid, 1310 La Hulpe, Belgium
Tel: +32-2-740 00 50 Fax: +32-2-740 00 59
email: info@eetimes.be.
VAT Registration: BE 461.357.437.
RPM: Nivelles.
It is is free to qualified engineers and managers involved in engineering decisions – see:
http://www.edn-europe.com/subscribe

CONTACTS

PUBLISHER
André Rousselot
+32 27400053
andre.rousselot@eetimes.be

EDITOR-IN-CHIEF
Graham Prophet
+44 7733 457432
edn-editor@eetimes.be

Patrick Mannion
Brand Director EDN Worldwide

CIRCULATION & FINANCE
Luc Desimpel
luc.desimpel@eetimes.be

ADVERTISING PRODUCTION & REPRINTS
Lydia Gijsegom
lydia.gijsegom@eetimes.be

ART MANAGER
Jean-Paul Speliers

ACCOUNTING
Ricardo Pinto Ferreira

european business press

SALES CONTACTS

Europe

Daniel Cardon
France, Spain, Portugal
+33 688 27 06 35
cardon.d@gmail.com

Nadia Liefsoens
Belgium
+32-11-224 397
n.liefsoens@fivemedia.be

Nick Walker
UK, Ireland, Israel, The Netherlands
+44 (0) 1442 864191
nickjwalker@btinternet.com

Victoria & Norbert Hufmann
Germany, Austria, Eastern Europe
+49 911 93 97 64 42
sales@hufmann.info

Monika Ailinger
Switzerland
+41-41-850 4424
m.ailinger@marcomedia.ch

Andrea Rancati
Italy
+39-02-284 6716
info@silvera.it

Colm Barry & Jeff Draycott
Scandinavia
+46-40-41 41 78
jeff.draycott@womp-int.com
colm.barry@telia.com

USA & Canada

Todd A. Bria
West
+1 831 477 2075
tbria@globalmediasales.com

Jim Lees
PA, NJ & NY
+1-610-626 0540
jim@leesmedia.com

Steve Priessman
East, Midwest, South Central & Canada
+1-630-420 8744
steve@stevenpriessman.com

Lesley Harmoning
East, Midwest, South Central & Canada
+1-218.686.6438
lesleyharmoning@gmail.com

Asia

Keita Sato
Japan
+81-3-6824-9386
MIshida@mx.itmedia.co.jp

Grace Wu
Asian Sources Publications Asia
(886-2) 2712-6877
wug@globalsources.com

John Ng
Asian Sources Publications Asia
(86-755) 8828 – 2656
jng@globalsources.com