boards & solutions + ECE

Combined Print Magazine for the European Embedded Market

March 01/17

Quick IoT Connectivity

from Specialised Development Kits

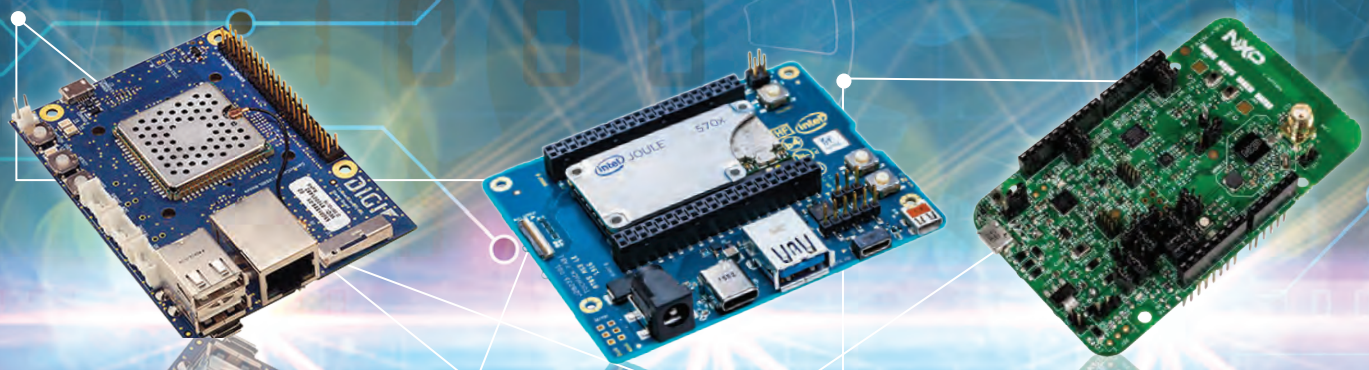

SPECIAL ISSUE

embeddedworld2017 Exhibition&Conference ...it's a smarter world

Dear Readers,

When we look back 15 years there was a small exhibition hall (the garage) in Sindelfingen in which the first predecessor of Embedded World was organized. Only a handful of exhibitors presented their offerings to the manageable number of visitors. Bad whether conditions did impair the start further. But despite all theses starting issues – today developed into a global market leader in the exhibition business.

This year the international embedded community once again will gather from 14 to 16 March 2017 in the Nuremberg Exhibition Centre. There will be the in the meantime biggest trade fair for embedded technologies taking place for the 15th year running. Trade visitors and conference delegates can already look forward to three exciting days at an event that is according to the organizer NürnbergMesse unrivalled in its compactness and offers direct contact with the most important trends and trendsetters. The success story of the embedded world exhibition & conference continues: already, the exhibition space reserved is higher than the previous year and is set to increase even more. For the first time, 1,000 exhibitors in six halls are expected to take part in the event. These figures once again highlight the fact that this exhibition on all aspects of embedded systems is the No. 1 international platform for the sector.

Two top-class congresses, the embedded world conference and the electronic displays conference, will take place alongside the exhibition. The embedded world conference covers the entire range of hardware, software and tools for the development of embedded systems. In 2017, one of the main thematic focus areas will be "Securely connecting the embedded world". The success of the embedded world conference is rooted in the direct involvement of the entire embedded community. The conference contents, carefully put together by an international jury, guarantee the necessary knowledge transfer and the annual exploration of the latest developments and trends.

The electronic displays conference has become established as the most important European B2B platform for display technologies. The topics discussed range from current display technologies like LCD, OLED or ePaper through flexible and 3D displays to GUIs and touch screens. In addition, the conference will also look at the wide range of applications for electronic displays and the development of the market.

Though, it is no surprise that this special issue of ECE/B&S one of the leading magazines for the embedded community covers all aspects of this industry as well. Enjoy reading!

And if you like to contact the team of ICC Media in Nuremberg there is the chance at booth 524 in hall 3A. We are looking forward to your visit.

Yours Sincerely

Wolfgang Patelay
Editor

Cover Photo:
Mouser Electronics

Cover Story:

Quick IoT connectivity from specialized development kits

Page 6

Specifically designed development kits provide hardware, software, firmware and tools for the IoT system, letting the developer concentrate entirely on the application development for his unique implementation. With such a wide variety of development kits available, getting an IoT system up and running has never been easier.

API for flash controller opens realm of custom applications

Page 12

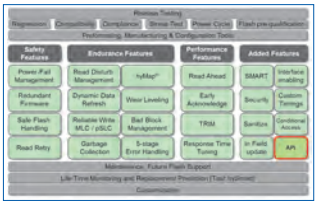

This article discusses the many applications accessible through Hyperstone Application Programming Interface (API). Users now have the option to supplement dedicated NAND flash management firmware from Hyperstone with their own Customer Firmware Extension (CFE) developed with the API.

A quick and easy way to design sensorless FOC motor controls

Page 22

This article describes easy entry for designers into the complex world of sophisticated motor control systems through either of two flexible, scalable, easy-to-use and affordable platforms from Infineon, catering for companies with different levels of existing competence.

Box PC with SMARC 2.0 module based on the Intel Apollo Lake processors

Page 25

This article introduces an all-new Box PC developed as an IoT gateway. Not only is the system standard eNUC for Box PCs new, the SMARC 2.0 module standard is new and the Intel Atom, Celeron and Pentium processors that are being deployed in the system were just launched in October 2016.

System design enablement and verification for the IoT

Page 36

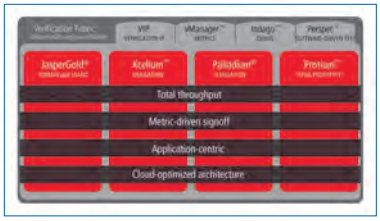

This article introduces some new additions to the Cadence Verification Suite, which makes it the go-to solution for verification and software development for designs in the new age of the Internet of Things.

Quick IoT connectivity from specialized development kits

By Mark Patrick, Mouser Electronics

Specifically designed development kits provide hardware, software, firmware and tools for the IoT system, letting the developer concentrate entirely on the application development for his unique implementation. With such a wide variety of development kits available, getting an IoT system up and running has never been easier.

■ The Internet of Things (IoT) is the main topic of conversation throughout the industry at the moment. Some experts will tell you that it is the next big thing, while others will tell you that it is already here. Whatever side of the discussion you are on, it is important to understand the concept behind the IoT and how it is supposed to benefit designs.

In its purest of forms, the IoT is all about connecting things (non-human) together to make decisions as quickly and accurately as possible. The things can come in many different forms – from smart objects, such as phones and watches, to small standalone nodes, the size of a postage stamp, but in general they all have common components and traits. These things are part of a larger system, and they generally acquire data from the environment through sensors, the data is then processed by a microcontroller and the processed data is sent back to a node to be passed on through the Internet to a server, which then may make a decision to take action on the data, or store it for future use.

Designing products for the IoT is much the same as designing microcontroller-based products for any other application; usually the first step is choosing the correct development kit, and to do that requires a thorough understanding of the target application. For IoT applications there could be three different types of projects. Cloud Connected: these projects can be the largest of all. They have the potential to connect systems from around the world to a cloud infrastructure. The cloud is made up of servers using dedicated software to analyze and process the data collected from the systems. Local Networked: these systems are generally a set of systems tied to a single location and may use a local intranet to communicate. Gateways: gateway systems are often used to retrofit existing systems to the Internet.

Developing products for cloud-connected systems highlights the main differences between developing for the IoT and developing a traditional embedded system. As mentioned already, the cloud consists of one or more servers to provide storage and services for the overall system. The nodes that make up the things of the IoT connect to these cloud servers. While the nodes gather the data, the cloud stores the information to build a historical picture of the operation of the system, and analyzes and processes the data in real-time to tailor the system to its optimal efficiency.
Trying to build a cloud connected IoT system from the ground up has the potential to be a daunting task. It could be expensive and time consuming and require expertise in many different engineering disciplines. The servers need to be set up with hardware and software. Drivers and application code would also be required for both the server and the IoT node. The hardware would need to be designed for the node. After this, the whole system would be required to be tested and debugged to ensure the connectivity between the servers and the node was seamless. All the tasks described would need to be completed before the application code for the system was written. Luckily for developers, there are cloud development tools available that can provide an ecosystem with various connectivity options and application code that already has been tried and tested. These tools can allow the developer to get a prototype up and running in a short amount of time – sometimes under a day.

To give an example of the different options that one manufacturer can offer developers, Intel has two boards designed to facilitate cloud development, or even be used as end devices. The Intel Edison measures only 25 mm x 35.5mm, enabling it to be integrated into products that need to be situated in small areas. Although small, the board packs a powerful punch through its 22nm Intel Atom SoC with a dual-core, dual-threaded CPU running at 500MHz, as well as a 32-bit Quark MCU running at 100MHz. The processors are supported by ample memory of 4GByte flash

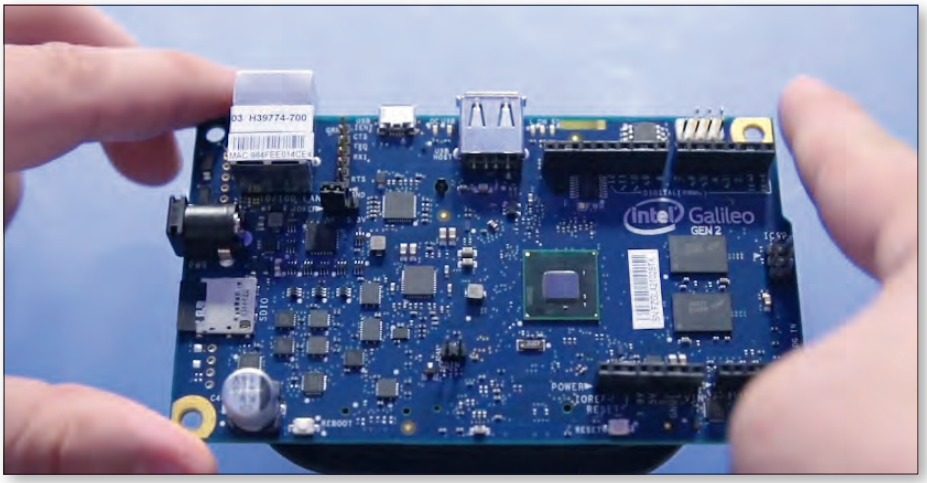

Figure 1. The Intel Galileo development board

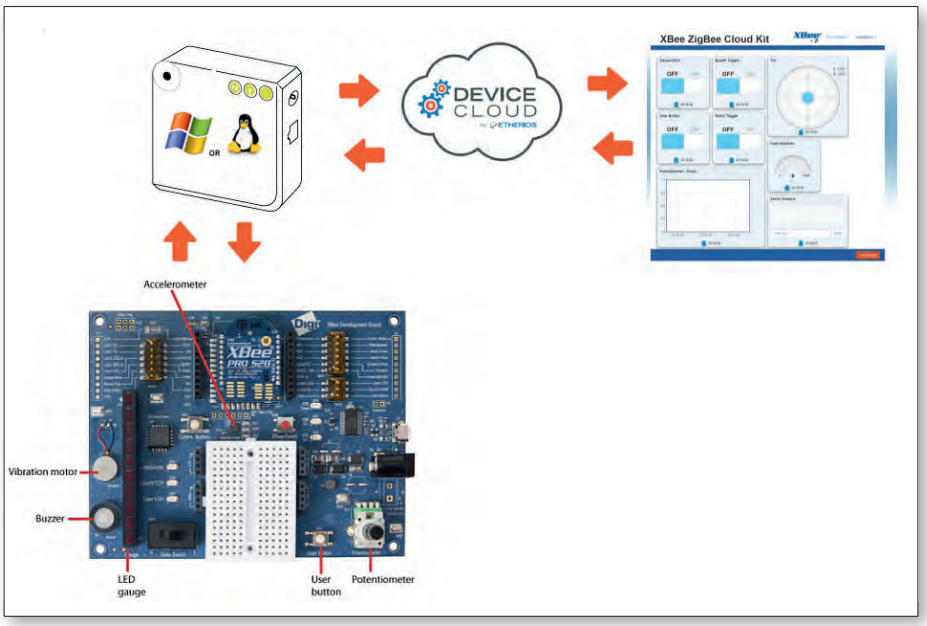

Figure 2. The XBee ZigBee Cloud Kit translates the ZigBee data into TCP/IP packets and forwards it to the Cloud via Ethernet or WiFi.

and 1GByte RAM. To provide the connectivity required for IoT applications, WiFi and Bluetooth are included. To support the board and make the server aspect of the system easier to integrate, Intel offers Intel IoT Analytics Cloud Server. This allows to collate, store and analyze data from the Edison without having to develop a custom server solution. The cloud server also acts as a development tool. To support application development and debug on the Edison, Intel provides an Arduino development environment.

The second Intel IoT board is larger and offers more features over the Edison. The Intel Galileo Gen 2 Arduino-certified board has a 400MHz 32-bit Intel Quark X1000 SoC at its heart and offers a raft of common embedded connectivity options, including SPI and UART. Connection to the Internet is taken care of with 10/100 Ethernet. As with the Edison board, the Galileo 2 is supported in the cloud by IoT Analytics Cloud Server. Both the Edison and the Galileo 2 development kits allow the gathering and transportation of data, with easy integration to cloud services. They are basic kits that offer easy initial migration to IoT development, but in some cases a more comprehensive solution with more powerful features may be required.

Microchip take the IoT concept a stage further with a cloud development platform that uses Amazon Elastic Compute Cloud (EC2) infrastructure. EC2 allows communication in both directions, both from the node to the cloud and back from the cloud to the node. This ability allows the Microchip development platform to be controlled directly from the cloud without any human intervention whatsoever. The main difference between developing for cloud systems and locally networked IoT implementations is that the local network may not necessarily require a TCP/IP communications link like Ethernet or WiFi. Bluetooth and Zigbee are two very popular standards which can be used for this type of networking and have other useful features such as lower power consumption and more resistance to industrial interference. The Texas Instruments (TI) CC2650STK SimpleLink Bluetooth Smart SensorTag IoT Kit is a good example of a development kit intended for a local network IoT application. The kit allows the developer to gather information as well as process the data and send it on. To facilitate the gathering, the board includes ten sensors for different applications including acceleration, temperature and ambient light measurement. The kit can also use an Android or iOS app, which connects to the development kit via Bluetooth. The app is used to display the readings from the sensors and provide basic analytics. It can also be used to control the sensors.

Another kit that provides a similar set of capabilities is the Sensor Puck from Silicon Labs. This kit offers a wide variety of sensors for ambient light, UV index, ambient temperature and humidity. The Sensor Puck kit also boasts an app for iOS or Android, which is connected to the board through a Bluetooth Low Energy (BLE) link. A useful feature on the Sensor Puck board is support for BLE broadcast, which can allow the app to collect data from multiple Sensor Puck kits. The developer can then write an application that analyzes the data from multiple sources.

Gateway devices allow data to be sent to the cloud from devices that are not able to support TCP/IP connections. These devices communicate using wireless standards like ZigBee instead of WiFi or Ethernet. Because ZigBee isn't based on TCP/IP, it doesn't assign IP addresses. Therefore, if the developer wants to send the data over the Internet without the extensive time and money penalties that bespoke development would bring, a gateway has to gather the information from the ZigBee network and translate it. For these types of applications, Digi has launched its XBee ZigBee Cloud Kit.

The Cloud Kit translates the ZigBee data into TCP/IP packets and forwards it to the cloud via Ethernet or WiFi. Digi has a sponsored device cloud account to host the data uploaded from the gateway. The developer can then use a browser to access the cloud control panel and interact with the data through the XBee Cloud Kit. Intel also has a solution for gateways, intended for IoT designs that are more complex. As well as supporting 10/100/1000 Ethernet and WiFi, the DK300 Series Gateway Solution also can send data to the Internet through cellular 2G/3G/4G. To gather the initial data from the nodes, the gateway boasts

Bluetooth, USB, serial ports and ZigBee. A developer starting from the ground up will have to consider the cost and time to design the complete system, as well as the numerous different types of specialist skills that are required to implement such a system. The company would have to have access to personnel with expertise in wired and wireless communications, embedded hardware and software design, and driver development for the data gathering nodes, as well as the skills such as cloud and mobile app development for the IT side of the system. ■

Product News

Hall-Stand 4868

Microchip: online store now includes all former Atmel products

Microchip announced that all former Atmel products are now available on microchipDIRECT, Microchip's full service channel. For the first time, customers can purchase products such as AVR and SAM microcontrollers and development tools direct from the manufacture.

News ID 4868

Hall-Stand 4919

SE Spezial-Electronic: Europe-wide 24-hour programming service for Si-Time's MEMS oscillators

As the first authorized distributor in Europe SE Spezial-Electronic offers as of now for all standard MEMS oscillator families by Si-Time an in-house 24-hour programming service. The selection of the respective components including definition of the desired technical components occurs simply by mouse click in the newly created e-shop. The components programmed inhouse by the service center will be delivered 24 hours after receipt of order at the latest. MEMS-based oscillators are no longer a niche product. Due to the manifold technical and economic advantages, they conquer more and more of the markets so far dominated by quartz-based components.

News ID 4919

Hall-Stand 4863

Maxim: remote tuner solution supporting worldwide radio standards for automotive

Using the remote tuner solution from Maxim Integrated Products, designers can significantly simplify the head unit design of a vehicle and reduce cables. The MAX2175 RF to Bits tuner within the solution eliminates the need to rework the vehicle's hardware to support worldwide radio standards, allowing updates by simply changing the vehicle's software.

News ID 4863

Hall-Stand 4855

Microchip: MOST150 technology ensures highest quality in-vehicle infotainment

Microchip announced that Audi has included MOST150 technology in their latest compact SUV model, the Q2. Media Oriented Systems Transport (MOST) technology is designed for in-vehicle infotainment platforms. MOST enables the deployment of a large variety of automotive premium surround sound audio systems and driver information systems such as full digital instrument clusters and heads-up displays. Audi has been using MOST150 technology for many years, also deploying it in their best-selling A4 Sedans, the Q7 SUV and TT Coupe models. This Q2 deployment utilizes the MOST150 Intelligent Network Interface Controllers (INIC) OS81110 and OS81118.

News ID 4855

Hall-Stand 4960

Mouser: 32-bit prototypes out of the box with PIC32MX470 Curiosity Dev Board

Mouser Electronics is now shipping the PIC32MX470 Curiosity development board from Microchip Technology. The cost-effective, fully integrated development platform enables rapid prototyping of 32-bit designs using the onboard Microchip PIC32MX470 microcontroller.

News ID 4960

Hall-Stand 4873

Toshiba: e·MMC ver.5.1 compliant Embedded NAND Flash memory products

Toshiba announced the launch of JEDEC e·MMC Version 5.1 compliant embedded NAND flash memory products with an enhanced operational temperature range of -40 to +105°C. The new products integrate NAND chips fabricated with 15nm process technology and are designed for industrial applications, including PLC, CoMs and factory automation equipment, and can also be used in a wide range of consumer applications. The line-up offers densities of 8GB, 16GB, 32GB and 64GB.

News ID 4873

Hall-Stand 4949

Renesas: dedicated motor control circuit technology for automotive MCUs

To increase the range of motor-driven vehicles, it is necessary to boost the energy efficiency of motor control. To achieve this, it is important not only to make mechanical improvements to the motors themselves, but also to enhance the functions and performance of the electronic control units that control the motors. ECUs capable of supporting next-generation EVs, HEVs, and PHEVs require advanced functionality and complex control software, and this results in a substantial increase in the operation processing load placed on the MCUs used in these ECUs.

News ID 4949

Hall-Stand 4888

Mouser debuts simple navigation of product datasheets

Mouser Electronics is constantly improving and simplifying how customers can browse and search for products on its website, where customers now can find product information more easily via Datasheets, Images and Newest Products, in addition to searching for products through the site's parametric Product Search.

News ID 4888

Hall-Stand 4923

Arrow: the journey from design to production at Embedded World 2017

At Embedded World, under the heading 'Think. Create. Produce,' Arrow Electronics is showcasing a range of products and services that can assist members of the maker community in transforming their innovative ideas into production. Start-ups and established companies can all benefit from a suite of online tools that, combined with Arrow's breadth of suppliers and global logistical capabilities, enable time-to-market to be shortened. Arrow will also be introducing three new products in the 96Boards specification at the show. Visitors can learn how to benefit from Arrow's relationship with crowdfunding pioneer, Indiegogo. Working with Indiegogo, Arrow is able to put critical resources in the hands of entrepreneurs, such as components procurement and online design tools, and also to share its expertise in the journey from design to production. There will also be the possibility to explore the comprehensive mixture of design tools, online engineering collaboration, reference designs and how-to content that Arrow is now making available via its website. Arrow will also demonstrate a free, integrated, cloud-based version of the Cadence OrCAD Capture design solution that makes it easier for design engineers and makers to integrate component research and selection within their design environment. OrCAD Capture Cloud saves significant development time, helping entrepreneurs get their products to market more quickly and cost effectively. With the Internet of Things becoming an essential element of so many embedded projects, Arrow has developed a range of IoT solutions from individual components through modules, system on module units, and more sophisticated proofs of concept. The Embedded World booth will feature a range of community boards together with Arrow's own development modules including the SmartEverything IoT board, a prototyping platform for M2M applications, and a new entry-level variant of the ARIS IoT hardware and software platform.

News ID 4923

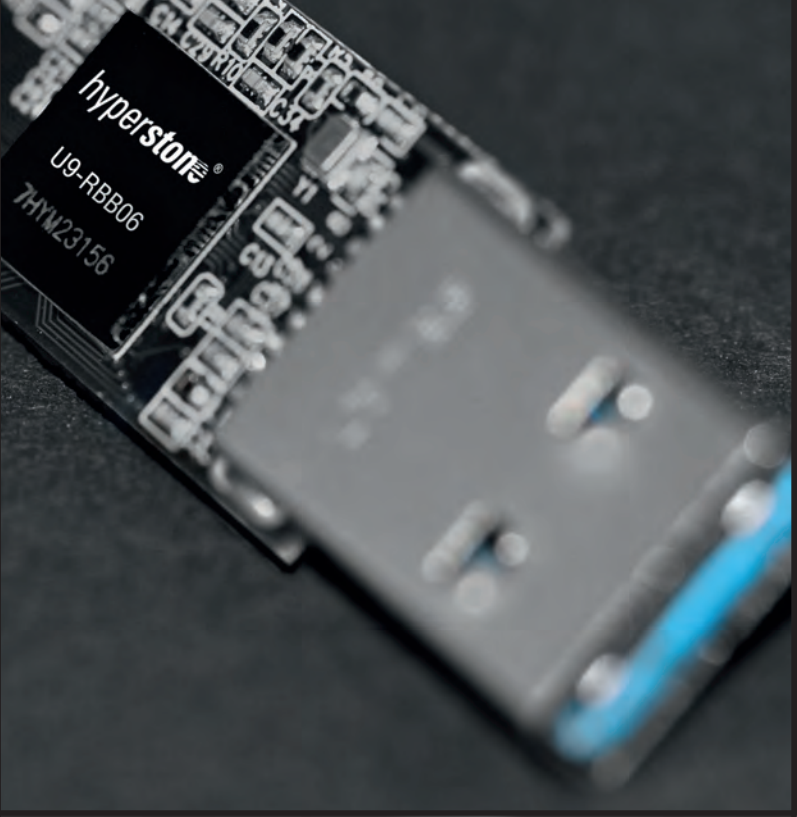

API

Security | IoT | Wireless | Your Solution

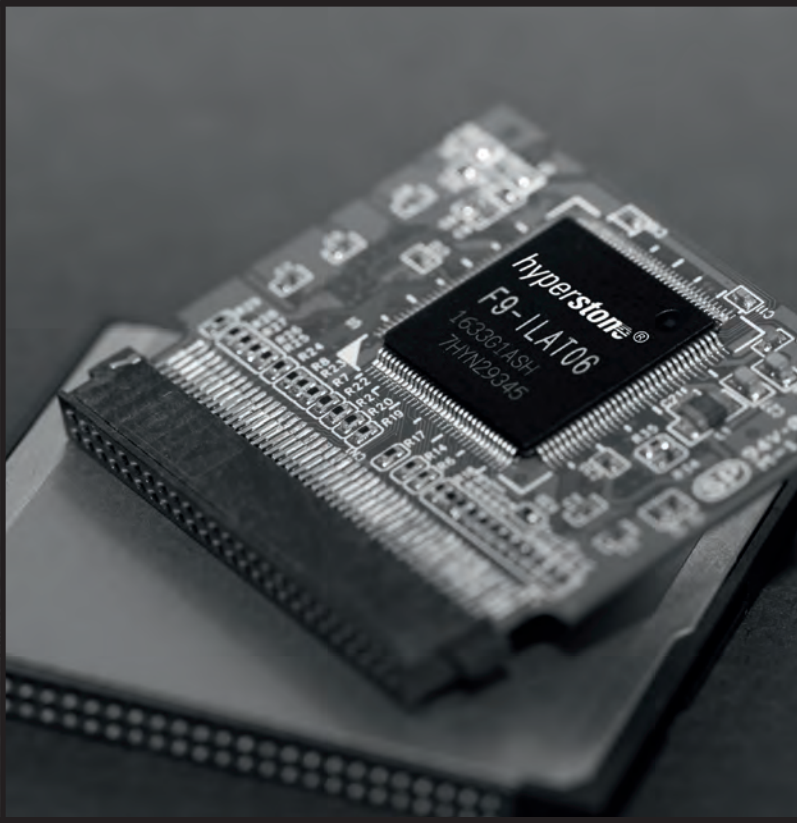

Our API – Your Individual Application

- Flash Memory Controllers for industrial/embedded applications
- Develop your own application with Custom Firmware Extension
- Full control and ownership of application IPs and releases
- Easily integrate key differentiators
- hyReliability™ Flash Management to increase lifetime

hyperstone®

API for flash controller opens realm of custom applications

By Damien Col, Hyperstone

This article discusses the many applications accessible through Hyperstone Application Programming Interface (API). Users now have the option to supplement dedicated NAND flash management firmware from Hyperstone with their own Customer Firmware Extension (CFE) developed with the API.

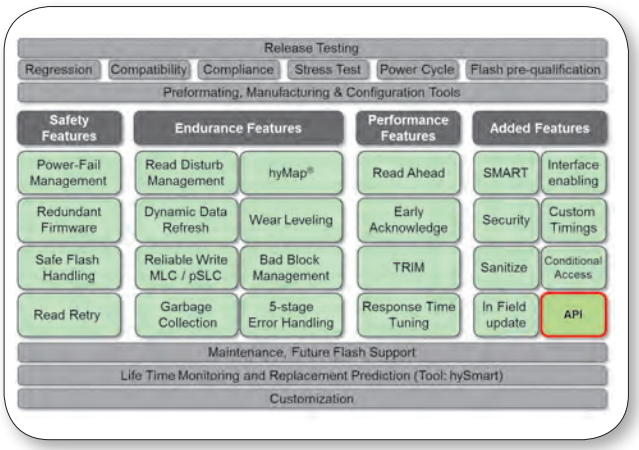

Figure 1. hyReliability flash management with API enhancement

■ Hyperstone have earned the reputation of providing the market with the most robust embedded storage solution by implementing proprietary Intellectual Property (IP) and patented techniques to manage the flash memory. Its flash management firmware is the cornerstone of its expertise, and the guarantee for long-term support for new flashes. With the addition of an Application Programming Interface (API) layer, and its unique offer among all flash memory controllers, Hyperstone now becomes a system solution enabler.

The API increases the realm of possibilities, by allowing customers to use world class Hyperstone Flash Management firmware and develop a customer specific firmware build, whether it is to add security or safety features, or to add interface to sensors, or communication interfaces for example. The ultimate goal is to integrate key differentiators. The developing party fully controls the ownership of the Customer Firmware Extension (CFE) and therefore protects its own IP. The CFE is completely independent of the firmware, and as such, can be reused on new generations of flash memories but also on different controllers, an assurance for the same long-term support.

As described in figure 1, API is another brick added to an already extensive set of advanced features to enhance the safety, reliability, endurance and performance of the latest generation of NAND flash memories, including 3D-NAND, and to enable maintenance and diagnostic required in the most demanding applications.

The API provides access to embedded firmware features and additional interfaces (figure 2) through the host interface in order to implement customized application added-value features, enabling a level of differentiation driven solely by the user's own expertise in the specific market to be tackled. The customer retains full control of his added functional features. The code can be changed at individual convenience, and never needs to be disclosed to Hyperstone.

The CFE becomes part of the final firmware code and is integrated alongside flash management advanced features. A majority of them are executed in the background and are transparent to the user. Those are associated with the safety, the endurance and the performance of the flash memories. These are the results of Hyperstone expertise to bring out the best of each generation of flash memories (SLC, pSLC, MLC, TLC, 3D-NAND).

The user binary can be written or updated by simple data write commands (prefixed by a vendor command that unlocks access to the reserved address range). As such, the customer code is also maintained under the same flash data management procedures (Wear Levelling, Error Correction, etc). If the customer data is moved, it will be handled by the flash translation layer (FTL) which translates logical to physical accesses. By means of call back functions, customers can intercept for example read or write commands, and add own routines. Also for example, based on vendor commands, users can establish communication between any kind of application software running on the host system with either the controller basic firmware (e.g. SMART) or certain of its hardware resources (e.g. AES, SPI, I2C or ISO7816) as well as tunnelling commands through the controller to a connected device such as a smartIC via the ISO7816 UART or a WLAN baseband. As an example, the CFE can have access to a range of features of Hyperstone advanced flash firmware, allowing it to easily manage its own private data outside the commonly user-accessible address space of a storage device.

As the customer binary is not dependent on the flash type used, the API program indeed only needs to be developed once and the same program can be used also if the flash type or Hyperstone controller changes. This way the system can always run with the latest flash technology, but without any additional development effort. Hyperstone firmware and the

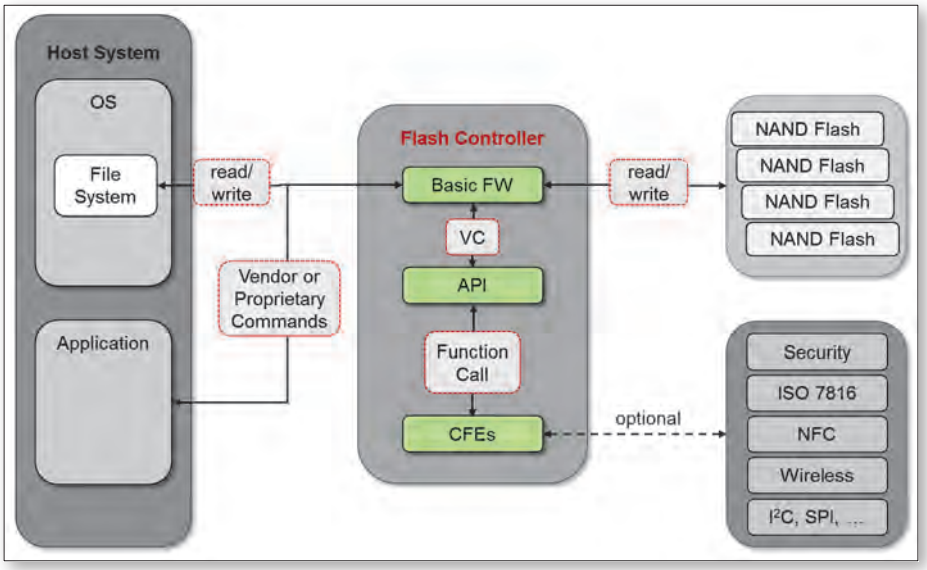

Figure 2. API hardware and firmware eco-system

CFE work in osmosis (figure 3). Hyperstone brings all its expertise in the flash management firmware, and will update it to support any new or changed technology, directly benefiting from its close relationship with the flash manufacturers. Complementarily, the CFE is the result of the customer expertise in its field of application. Both work hand-in-hand, but are completely independent from each other, hence providing a full flexibility. API serves several of our customers' needs: it enables new applications, strengthens them in their domain of expertise, creates differentiation features against one another, facilitates new feature development without being limited by Hyperstone resources, and enables new firmware feature development. It would be impossible to list the number of applications that can be addressed with the API. However, hot topics are certainly secure storage and secure transactions. For example, certain applications require a dedicated proprietary firmware in order to utilize the ISO 7816 or SPI interface, to implement encryption key management, and to communicate securely and reliably with a host application. This includes: secure storage, mobile payment, mobile Pay TV, secure mobile communication (data/voice), digital rights management (DRM) (conditional access, license management, content protection), data privacy (e.g. medical), FIDO U2F security key (anti-phishing protection/built-in smart card) with AES data encryption, authentication and IoT for wearables. In addition, other serial interfaces available on Hyperstone controllers can be enabled and made accessible for applications. Behind the scenes, the CFE is stored as binary code in a reserved range not presented to the host (figure 4). This is called the CFE load area.

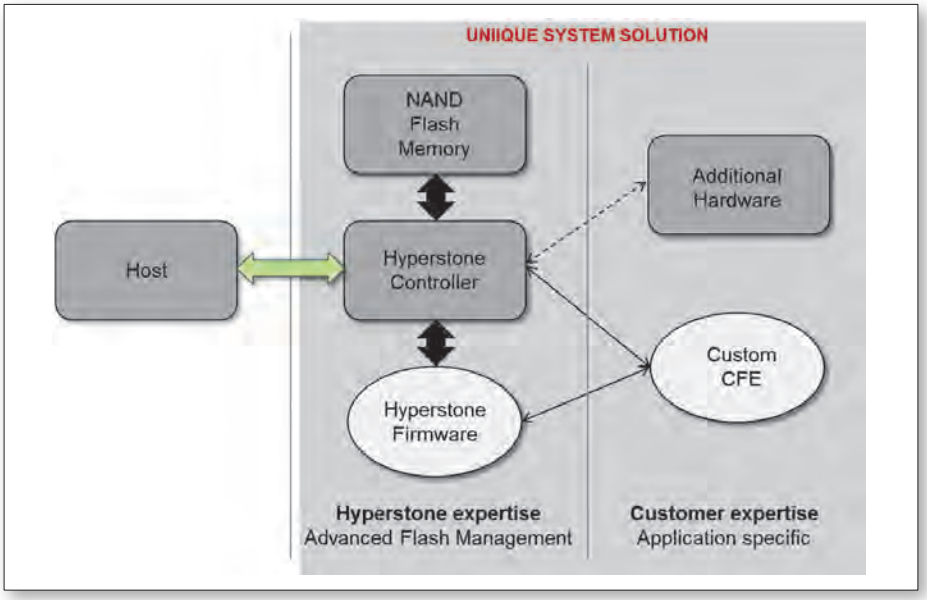

Figure 3. Hyperstone hardware and firmware + flash + CFE = unique system solution

The source code of such CFEs is solely owned and controlled by the customer. A customer firmware extension (plug-in) binary file, also referenced as CFE, consists of: an authentication section (always located at the beginning of the file), a header section (located directly after the authentication section), and up to three data sections (static, volatile, overlay). This is however transparent to most developers who will not have to worry about it. All firmware extension binaries are encrypted with AES-256 in CBC mode.

To execute API commands through vendor commands, it suffices to register the vendor command in the CFE (either as DataToHostCallBack or DataFromHostCallBack) and to execute the vendor command using the proprietary Hyperstone hsfmt tool, delivered with the kit. The tool is supplied as C source code and can be compiled for different operating systems. Different options exist for debugging, such as using the debug interface of Eclipse or the debug-host via the debug-UART interface.

For further information, specific application notes are available. Breakpoints can be set dynamically at runtime. A simple debug possibility is the capability to send (debug) strings with the debug-UART and print them on the host. The key advantages for the API user can be summarized in the following set of added values: full control over application IPs, full control of releases, no synchronization needed with Hyperstone firmware releases, no source code transfer required, unique differentiator through application IPs, and fast software development. The API kit is available for purchase and without recurrent licensing cost including API function support from Hyperstone and firmware field update (FFU).

Different models exist for the development of new applications. The API software development kit (SDK) can be offered to system integrators or end customers directly. It is not required to understand the complex flash management features to develop a CFE. The kit makes it possible for anyone with basic firmware knowledge to develop a new application and implement key differentiators. In some other cases, it is also possible to outsource the tasks to third party software development service providers.

With Hyperstone flash controller API, many applications are imaginable to enhance the value proposition of your system. Security applications can be explained in more detail as an example, one among many others. With API, security features embedded in the Hyperstone controller, like AES encryption engine, SHA engine, secure random number generator can be enabled and managed. It allows the implementation of arbitrary security features

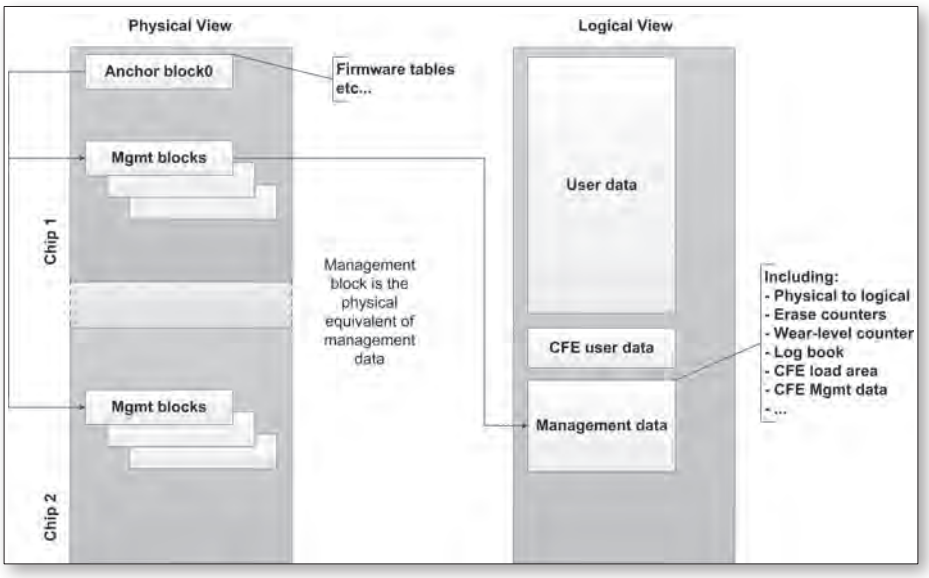

Figure 4. Anchor block and corresponding CFE load area (hyMap)

and access control systems. In this context, having full ownership of your system is vital. The realm of such applications is limitless, but to mention a few: build own access control system (key management/hashing), build and manage private data space, secure read/write, user key or automatic key generation, encrypted keys, and crypto erase, the fastest way to invalidate data. But the applications do not end here. As a further example, a wireless module can be connected to the flash controller, to get lifetime information, or to save or restore a flash memory image wirelessly.

Hyperstone flash controller API enables the development of dedicated proprietary features through a user-friendly environment, delivered in a kit. Using API, a customer firmware extension can be developed under the strict control of the system integrator, who keeps full ownership of the developed value-added application feature. The serenity of the CFE is also insured as it can be transferred to new generations of flash, without any additional investment. The porting effort across different Hyperstone controllers is minimal. The API is available for S8, U9 and subsequent designs. It has been used for applications now in mass production. Hyperstone will demonstrate the API at Embedded World (Hall 1/Stand 1-301). ■

Product News

Hall-Stand 4899

Atlantik Elektronik presents Bluetooth Smart SoC from Qualcomm

Atlantik Elektronik will present with the Bluetooth Smart 4.2 System on Chip CSR102x from Qualcomm Technologies a product family specifically designed to help engineers meet the needs of today's "always-on" world. The CSR102x family is optimized for specific applications in the Internet of Things, including wireless remote controls, simple smart watches, home automation solutions and beacons, where balancing performance, battery life and cost is critical.

News ID 4899

Hall-Stand 4898

Rutronik: Bluetooth Smart ICs with low current performance from Toshiba

Toshiba's three new ICs TC35678FSG, TC35678FXG and TS35679FSG, support Bluetooth Low Energy version 4.1 communications. They achieve class-leading low current consumption and are ideal for use in Bluetooth Smart devices, such as wearable technology, medical equipment, smartphone accessories, remote controls and a wide variety of emerging IoT solutions. The new ICs are available at distributor Rutronik as of now.

News ID 4898

Hall-Stand 4894

Mouser: DC-DC synchronous buck converters operate in 3.5 to 36V input voltage range

Mouser Electronics is now stocking MAX17242/MAX17243 step-down DC-DC converters from Maxim Integrated. These high-efficiency, synchronous converters with integrated MOSFETs help minimize external component count and reduce total cost of solutions for a variety of applications, including power regulation and general-purpose point-of-load designs.

News ID 4894

Security portfolio gives protection from industrial and IoT threats

By Stefan Ingenhaag, Renesas Electronics Europe

The industrial and IoT markets present security considerations unlike anything embedded developers have ever faced. By offering a comprehensive and robust set of security features in both hardware and software, Renesas Synergy Platform delivers a sure path to long-term product integrity.

■ With an estimated 14 billion devices connected to the Internet today and that number expected to jump to 50 billion by 2020, the Internet of Things (IoT) offers a tremendous market opportunity. But it also presents a major security risk. Why? First and foremost, the vast majority of all connected devices in use today feature inadequate security. A recent report from the Open Web Security Application Project (OWSAP) gives some insight into the scope of the problem. According to the report: 60% of devices with user interfaces are vulnerable to issues such as weak credentials, 70% of devices use unencrypted network services, 70% of devices, along with cloud and mobile applications, enable an attacker to identify valid user accounts through account enumeration, and 80% of all devices, along with cloud and mobile applications, fail to enforce a policy requiring passwords of sufficient length and complexity.

As more devices and systems connect to the network, the potential risk and implications of a security breach continue to climb. A Denial of Service (DoS) or Distributed Denial of Service (DDoS) attack on an industrial plant could result in a system overload, leading to unavailability of important services and notifications at a critical time, for instance an industrial cooling system that provides critical temperature notifications could be severely compromised if the notifications to system operator or commands from the operator are blocked due to DoS/DDoS attack(s) on the connected temperature monitoring and control system. Similarly a DoS/DDoS attack can interfere with road traffic control systems resulting in massive gridlock in the area it serves. From a national security perspective the implications are even more frightening. The National Security Administration (NSA) admitted in 2014 that the US is in a de-facto state of cyber warfare and that hackers had already infiltrated systems in reconnaissance missions. And senior military officers reported that a number of countries now have the capability to shut down the US power and financial services industries.

From an individual user standpoint, consumers are becoming more vulnerable as they increasingly rely on connected devices. The use of connected healthcare devices such as baby monitors, glucose readers and pacemakers leaves the health and well-being of consumers exposed to malicious add-ons that could render those devices useless or worse administer incorrect does of medication. As consumers increase the intelligence and interconnectivity of their homes, the risk of hackers stealing personal data increases exponentially. To get a sense of the scope of the problem, Renesas recently asked its customers what they see as the most dangerous threats in the IoT marketplace. They identified the following six types of threats. 1) An untrusted contract manufacturer cloning software or firmware or the security configuration of an MCU or product. 2) Hackers disrupting a product by replacing a genuine firmware during the installation phase with malware. 3) A hacker mounting an eavesdropping attack during firmware installation especially if security parameters are exchanged in the clear. 4) Privacy threats when system firmware is not physically protected and an attacker can then extract security parameters. 5) Attackers using an add-on program to damage or steal information. 6) Hackers taking advantage of a simple software update session to replace firmware with malware.

For the designers of Renesas Synergy Platform, one thing was immediately clear. To ensure devices built around this platform were truly safe and secure, they had to address all these threats. They realized this meant they had to build into the platform security features that would provide protection at each stage of the product lifecycle.

To address potential threats in virtually any application, Renesas engineers developed an integrated hardware/software platform that offers an unprecedented line-up of security

Unique ID	True Random Number Gen	Crypto HASH Functions	Symmetric Key Crypto	Asynmetric Key Crypto	Secure Key Storage	Read-Out Protection

Security Software Library

Threat	S7	S5	S3	S1
Product cloning	Best	Best	Better	Good
Product disruption with malware injection during update	Best	Best	Better	Good
Eaves-dropping during update	Best	Best	Better	Good
Privacy threat by firmware/data exposure	Best	Best	Best	Good
Add-on program to damage or steal	Best	Best	Best	Limited

RENESAS SYNERGY SECURITY PROTECTION

Figure 1. Renesas Synergy Platform offers a broad range of security features.

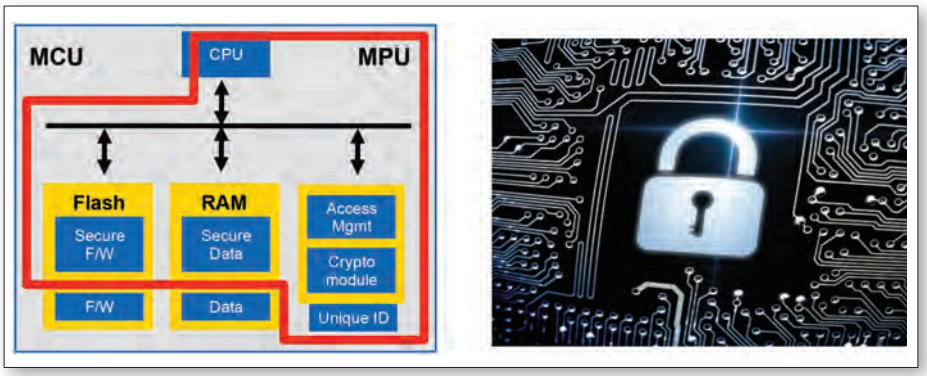

Figure 2. New features like key injection service and certifiable root of trust offer a solid defense against product cloning.

capabilities. Many of these new functions are implemented in hardware where they are less susceptible to attack. As an example, when each MCU in the Renesas Synergy family is manufactured, it is assigned a unique 128-bit ID which can be used to generate keys to protect applications and assist provisioning. Providing a significant improvement over pseudo-random generators, the Synergy random number generator meets the latest NIST SP 800-90 specifications and is tightly integrated into the cryptographic accelerators and key generator of the MCUs.

All MCUs in the Synergy product family, including the S1, S3, S5 and S7 Series controllers, feature accelerators for symmetric cryptography, HASH, a true random number generator and the ability to limit JTAG access. At the higher performance end of the product line, the S7 and S5 Series add accelerators for asymmetric cryptography, asymmetric key generation and key secure storage. To ensure that certain areas of memory can be accessed only by those with the right privileges, Secure Memory Protection Units (MPUs) in the MCUs are used. The platform plays a key role in ensuring the integrity and availability of systems. As an example, the separation of stakeholders is crucial to the integrity of many connected systems. By separating session keys and user data from equipment data and allowing the system to create a separate sandbox for configuration data, the Secure MPU, available on the Synergy MCUs, allows users to operate multiple stakeholders at the same time.

The risk of IP theft begins in the manufacturing phase as soon as the MCU leaves the hands of the design team. Say your design team is building a glucose reader. A less-than-honest contract manufacturer can reverse engineer a copyrighted algorithm to create counterfeit glucose reader MCUs to sell to competitors. Product cloning like this threatens to undermine sales revenue. It can also damage the original manufacturer reputation when lower quality devices hit the market and create support issues for partners who unknowingly purchased cloned parts. To protect against this threat, Synergy Platform developers integrated two innovative features into each MCU to help authenticate each device and minimize the threat of cloning. In the earliest stages of the manufacturing process, developers using the Synergy MCUs have the option of choosing Renesas or a third party key injection service. Key injection essentially assigns a unique identifier or birth certificate to each MCU. Early key injection allows code to be locked to a device. From that point forward a checksum of the code must achieve a known value to be accepted.

The earlier in the manufacturing process this unique identifier is assigned, the more difficult it is for anyone to steal the MCU identity. Today the Synergy S7 Series is one of a select few controllers offering key storage and generation on the MCU. The second feature Renesas is offering on the Synergy Platform to address product cloning is a certifiable root of trust. The root of trust serves as the basic foundation for security upon which other security components are built. It includes the key components in an embedded system that the operating system must trust and which must operate immediately out of reset including secure firmware, data, access management, the cryptography module, and the unique ID of the device.

The root of trust serves three key functions. First, it must measure and verify the software boot chain. This requires that it exist separately and underneath any system boot chain. Second, the root of trust must protect the cryptographic keys. This means that the root keys must be provisioned early and securely. The Synergy Platform early key injection, secure storage and limited JTAG access capabilities ensure it meets these requirements. Third, the root of trust must perform device authentication. The Synergy platform performs this task via asymmetric key generation, the asymmetric crypto-accelerator, the true random number generator, and the symmetric crypto engine.

Once an IoT device is deployed to the field, every remote patch, software update or lifecycle maintenance routine poses a potential threat. Each of these functions must be performed remotely and securely. To protect against malicious overwrites, the Synergy Platform offers an authenticated boot capability that places code and its key in a secured flash area in the on-chip memory of the MCU. In figure 3, the traditional MCU on the left has simply allocated the customer authentication code to user flash. The code is therefore

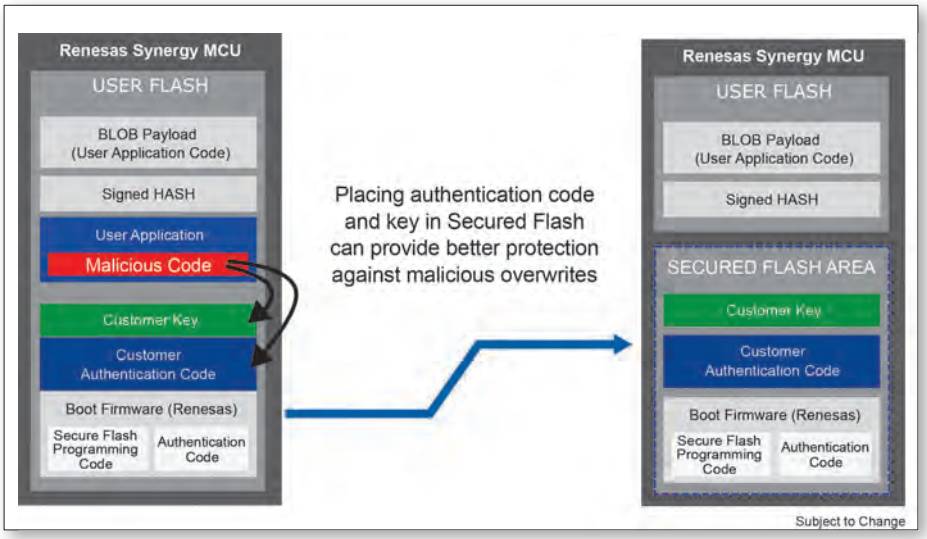

Figure 3. By placing authentication code and key in secured flash, Synergy's authenticated boot function extends product life.

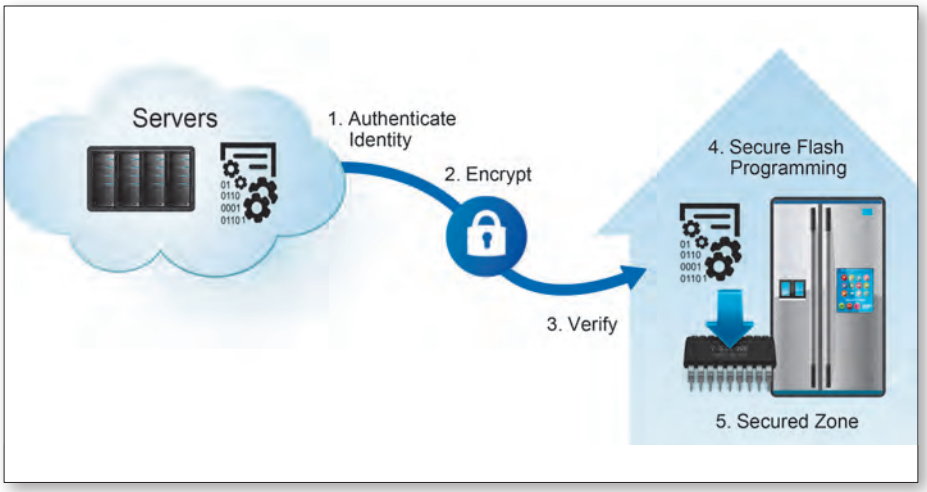

Figure 4. Synergy Platform offers a broad range of security features.

under threat of being overwritten. The Synergy Platform counter measure on the right side of the figure makes it far more difficult to hack by restricting JTAG access and protecting the authentication code in a secured flash zone of the Memory Protection Unit. The rest of the user flash security load remains unchanged. In this example, the BLOB payload would still likely be an update with the signed HASH authenticating that it comes from a trusted source and the expected size (HASH) to ensure nothing has been added.

Synergy authenticated boot feature can extend product life by ensuring a device continues to be protected in the event of a susceptibility. Once susceptibility is identified, authenticated boot will only allow the correctly identified owner to perform devices updates, it will only allow updates on signed and verified code from the server, and it will inhibit roll-backs of software to previous versions. In a post-deployment environment, users of connected devices or systems are constantly facing hacking threats. Hackers breaching networked smart homes can steal private data and undermine the operational capability of equipment on the network. Similarly, hackers attacking smart factories can slow or halt production or create a catastrophic failure.

To protect against device hacking, the Synergy Platform adds three key security features: trusted libraries of code, sandboxing and the ThreadX RTOS. Trusted libraries of code avoid buffer overflow and code injection by constraining inputs and communication. As an example, cryptography and secure boot components enable only a subset of inputs which helps create contagion control. Sandboxing is a feature which enables the Synergy Platform to maintain availability even while under attack. It allows the system to separate different aspects of the operating system and an application. Basically it functions by using an integrated MPU to segregate the on-chip memory into areas only accessible by privileged applications and areas which are open to communication. This segmentation allows users to create supervisory and user modes and ensures that any problems in user mode do not reach supervisory mode. In addition, with sandboxing users can create areas of memory as unreadable or unprintable. Once that division is set up, when an application goes out of bounds, as in a denial of service attack, the MCU will raise a flag in hardware to execute a handle to reset the system. That, in turn, will defend against the attack. Finally, the use of an industry-proven RTOS, Express Logic ThreadX, offers significant security advantages. As an operating system that has been used for years by tens of thousands of users, ThreadX offers users the reassurance that it has been developed, checked and certified to robust security standards.

At the development and manufacturing stages of the product lifecycle, embedded systems are particularly vulnerable through their firmware. One of the primary advantages of a connected device is the ability to update the firmware remotely. But this capability also creates vulnerabilities. In a typical example MCUs in a car assembly line programmed via firmware are injected with malware that undermines performance with potentially life-threatening consequences. Synergy Platform protects against this type of attack with a five-step authenticated firm-ware management program. This process begins by verifying the digital certificate and identity of the firmware update service. The Synergy Platform authenticates its communication channel by using several technologies including the unique ID, key secure storage, asymmetric cryptography, and the asymmetric key generator. Next, to protect the Binary from tampering and interception during the download, Synergy uses a Transport Layer Security (TLS) connection which employs the True Random Number generator and the Secure Crypto Engine with HASH.

Key to the platform ability to deliver high performance cryptography, the Secure Crypto Engine integrates a number of crucial functions including the True Random Number Generator, a symmetric cryptography accelerator, an asymmetric cryptography accelerator, a HASH accelerator, a stream cipher accelerator, and Secure Key Storage. Since the True Random Number Generator resides within the Secure Crypto Engine module close to the root of trust, is not software-enabled, and supports a smaller session key, it is less susceptible to attack. In step three the Synergy Platform verifies that the downloaded Binary is the same as the source Binary on the server by using its integrated HASH technology. Next, the platform only programs the authenticated Binary using the Secure Crypto Engine. Finally, to protect the binary on flash, the system creates a secured zone using the secure MPU and limited JTAG access. ■

Embedded/IoT Solutions

Connecting the Intelligent World from Devices to the Cloud

Long Life Cycle · High-Efficiency · Compact Form Factor · High Performance · Global Services

IoT Gateway Solutions

E100-8Q, E200-9B

Compact Embedded Server Appliances

SYS-5028D-TN4T

Network, Security Appliances

SYS-5018D-FN8T (Front I/O)

High Performance / Short Depth Solutions

SYS-1028R-WMRT

Cold Storage

SYS-5018A-AR12L

High Performance / IPC Solutions

SYS-6018R-TD (Rear I/O)

- Low Power Intel® Quark™, Intel® Atom™ Intel®, Core™ processor families, and High Performance Intel® Xeon® processors
- Standard Form Factor and High Performance Motherboards
- Optimized Short-Depth Industrial Rackmount Platforms
- Energy Efficient Titanium - Gold Level Power Supplies
- Fully Optimized SuperServers Ready to Deploy Solutions
- Remote Management by IPMI or Intel® AMT
- Worldwide Service with Extended Product Life Cycle Support
- Optimized for Embedded/IoT Applications





Intel Inside®.
Powerful Productivity
Outside.

TITANIUM

We Power the Cloud™

Learn more at www.supermicro.com/embedded

USB Version 3.1 – a crucial turning point for industrial applications?

By Nicolas Prou, Würth Elektronik eiSos

Universal Serial Bus (USB) has a strong foothold in consumer electronics, but has yet to prove itself in industrial applications. Electronic component manufacturers predict a significant increase in the future and see the introduction of the USB 3.1 standard as a turning point, if certain quality standards of the components are met.

Figure 1. USB-3.1 connector, C-type: the end of trial and error on insertion

■ IEEE 1284 and PS/2 ports on PCs are just two of many interfaces that USB has completely replaced as the standard in the consumer field. In industrial applications, USB has also increasingly been applied in the past, for example to connect mobile devices, to read out measurement data, or to import and install software updates onto machines. But why is USB not widely used in industrial applications even more? Depending on the application, restrictions in the specification could be one reason. But one thing not to be underestimated is that USB has an image problem.

USB connectors are considered to be cheap accessories for undemanding applications, but not a robust industrial connector. And that is not entirely unjustified. Many products simply had quality deficits; the price pressure on universal connectors was extremely high. For this reason, the raw materials used were often not matched to the demands of the application, which led to problems with solderability and consequently to weaknesses in the locking mechanism and durability. Not only the materials played an important role, but rather the mechanical design of the connector: insertion and withdrawal forces could have a negative impact. In the worst case, the connector destroyed the socket on insertion or pulling out.

Another latent issue: on account of the symmetrical data transmission, USB should actually be very interference-resistant. In the case of the inductive interference effect (magnetic field), the twisting of the wires achieves compensation of the interference effect. Theoretically, as a result of symmetrization of the partial inductances, the interference influences compensate each other. Practically, this interference immunity can be compromised, which led to use for some applications being considered unsuitable. The reason often lies in the lack of symmetry of the inputs/outputs of the USB controller or simply in the poor quality of the modules. The interface design (receptacle, housing) is inadequate. Poor ground reduces the shield attenuation of the cable. Filters have poor ground reference. The USB cable is asymmetrical, poorly shielded and has inadequate ground connection. The cable deteriorates the signal quality, radiates signal harmonics and has insufficient shield attenuation towards external interference sources. Only those developers in industry who anticipated that the controller bought-in or the cable used by the end user may not be sufficiently immune to interference and looked into additional EMC protection circumvented this problem. Manufacturers and developers that ignored the EMC issue with a view to the theoretical properties of the USB connectors met with problems in practice.

Those who previously considered USB to be unsuitable for the mentioned reasons will rethink their decision in view of the additional advantages of USB 3.1. Doubling the data rate to 10Gbps alone is a powerful argument. The significantly increased supply power from 9W up to 100W is also extremely interesting for industrial applications. This enables additional applications. The power supply can now also be reversed – i.e. from the client to the host. Last but not least: USB 3.1 does away with the tedious poking to insert the USB connector the right way around: The C-type plug no longer has a direction in the USB 3.1 version. Not only the developers of consumer electronics will want to allow their users to benefit from this faster and simpler handling. 10 Gbps, 100W, simple handling – with these arguments a broad field of industrial applications opens up for USB 3.1. USB can now replace technologies complicated in their use - PCI Express or Thunderbolt, also with the cost argument of mass production, of course.

The conceivable deployment scenarios for USB really are universal. A single connector is the answer to many different requirements, such as data exchange, power supply, video content and networks, but also as a replacement for D-Sub and HDMI, as well as modular sockets. Examples for typical applications are external hard disks and servers, medical

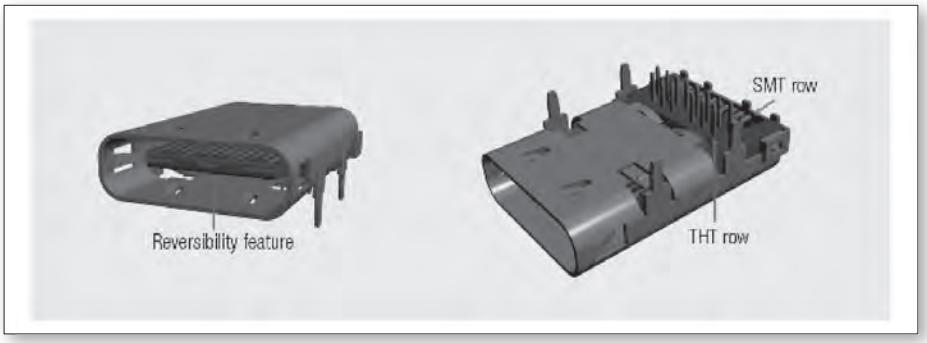

Figure 2. USB 3.1 C-type socket

A1	A2	A3	A4	A5	A6	A7	A8	A9	A10	A11	A12
GND	TX1+	TX1–	Vbus	CC1	D+	D–	SBU1	Vbus	RX2–	RX2+	GND
GND	RX1+	RX1–	Vbus	SBU2	D–	D+	CC2	Vbus	TX2–	TX2+	GND
B12	B11	B10	B9	B8	B7	B6	B5	B4	B3	B2	B1

Figure 3. Pin assignment USB 3.1 C-type socket

applications with integrated monitors, industry applications that require data backup, smart phones and all portable devices like sensors, power banks and small monitors. Nevertheless: even with the extended technical possibilities, the basic conditions for industrial use still need to be observed. For an industrial USB application, the issue of EMC protection outlined already is important and has to be credibly addressed in the applications. And those who do not want to risk the reliability of their application should not save too much on material and integrate high quality modules.

An example for one of the details that distinguishes the high-quality connector: the optimized connection angle. If the connection angle is greater than 40 degrees, the contacts could be damaged on inserting the connector. In theory, connectors have a chamfer of 30 degrees in order to limit the risk. In practice, this chamfer is not always integrated, however. So low cost connectors can damage the end application of the customer which leads to complaints, service costs, and the risk of suffering image damage - just because of an USB socket. The sockets on the market therefore have to be designed such that they are also suitable for non-compliant connectors. USB-3.1 modules with improved shielding and increased coating thickness of the contact surfaces to avoid high mechanical loads are the way forward - but are expensive, of course. But with a durability of 10,000 plug cycles, they are also suitable for devices that should last longer than short-lived consumer products.

What comes next with USB3.1? Firstly, there is a dynamic market development with consumer applications to be expected. Even today, many manufacturers of smartphones already use USB 3.1, C-type to supersede USB 2.0 or Micro USB 3.0, B-type. Connector producers expect industrial companies to follow suit within a year once development teams have evaluated their requirements and possible constraints on the use of the new technology. Especially when attention is paid to high-quality material and clean EMC solutions, USB 3.1 is true progress. ■

Figure 4. USB 3.1 C-type views

A quick and easy way to design sensorless FOC motor controls

By Ivan Dobes and Frederic Zimmer, Infineon Technologies

This article describes easy entry for designers into the complex world of sophisticated motor control systems through either of two flexible, scalable, easy-to-use and affordable platforms from Infineon, catering for companies with different levels of existing competence.

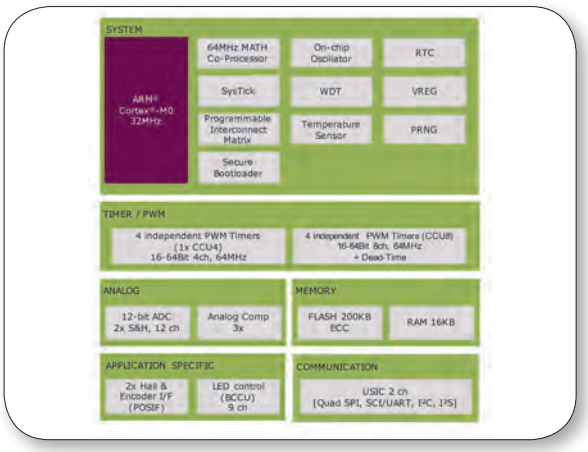

The XMC 1300 series with its dedicated peripherals is suited for cost-efficient, sensorless or sensor-assisted motor control

■ Motor control application designers are under pressure today from the markets, competitors and regulators to optimize the performance and efficiency of their systems, while keeping the costs of the electronics systems to the minimum. Applications ranging from pumps, fans and compressors in home appliances to drones and e-bikes require sophisticated control systems able to respond reliably to dynamic loads and rapid changes in motor speed while keeping the efficiency at the highest possible level, and at the same time meeting the other application requirements, such as low acoustic noise, for example. The ideal solution to this design challenge is the sensorless field-oriented control (FOC) of permanent magnet synchronous motors (PMSMs). However, designing such a system requires expertise in advanced motion-control algorithms, software optimization, and motor performance.

Infineon offers two types of controllers which are capable of running sensorless FOC control of PMSM motors: XMC MCUs, based on standard ARM-Cortex-M0 core, and iMOTION IRMCK099 motion controller IC. These two types of controllers serve different types of designers. XMC targets designers and companies with solid or extensive know-how in motor control and embedded programming, which can then use the motor control software that Infineon provides as a basis for their own motor control and system software. On the other hand, iMOTION IRMCK099 serves companies which do not have sufficient motor-control know-how, or their core competence is based at the system level. These companies can then use the IRMCK099 as a turnkey solution for motor control, while focusing on differentiation at the system level.

With four product series and more than 100 products, the XMC1000 family caters for a wide range of applications. The XMC1000 derivatives offer between 8 KB and 200 KB flash, efficient PWM timers, 12-bit A/D converters and flexible serial communication interfaces. For the simplest applications, the XMC1100 entry-level series offers a basic range of functions to facilitate entry into the XMC world. The XMC1200 series brings in additional application-specific features, such as a Brightness and Colour Control Unit (BCCU) for LED colour and brightness control as well as capacitive touch control and LED matrix control unit. One performance level higher, the XMC1300 microcontroller series adds interfaces and hardware blocks optimized for motor control. In addition to the high-performance Capture and Compare Unit (CCU8), it also incorporates a position interface (POSIF) for the precise detection of the motor position by external Hall sensors or encoders. A MATH coprocessor (XMC1302) accelerates key mathematical operations required for execution of complex motor control algorithms. These features of the XMC1300 series enable both sensor-assisted and sensorless field-oriented control of electric motors. The latest XMC1400 series offers a performance boost of up to 70% thanks to a faster CPU clock and built-in pre-fetch buffer. This series also features a CAN interface which is very often used as a communication protocol in the industrial environment. The entire XMC1000 family offers variants for the temperature range up to +105 °C.

For numerous industrial applications, and particularly for efficient motor controls, high-speed, precision A/D converters are required. The A/D converter of the XMC1000 family can be configured with resolutions of 8, 10 and 12 bit. This allows the conversion time and resolution to be optimised for the application. Beside high resolution and sampling rate, the precise setting of the time of a measurement is particularly important. The XMC1200, XMC1300 and XMC1400 series thus allow sampling times to be defined via programmable, hardware-controlled sequencer in accordance with the switching patterns. The MATH co-processor of the XMC1300 and XMC1400 series can be used for vector rotation (PARK transformation) with 24-bit resolution.

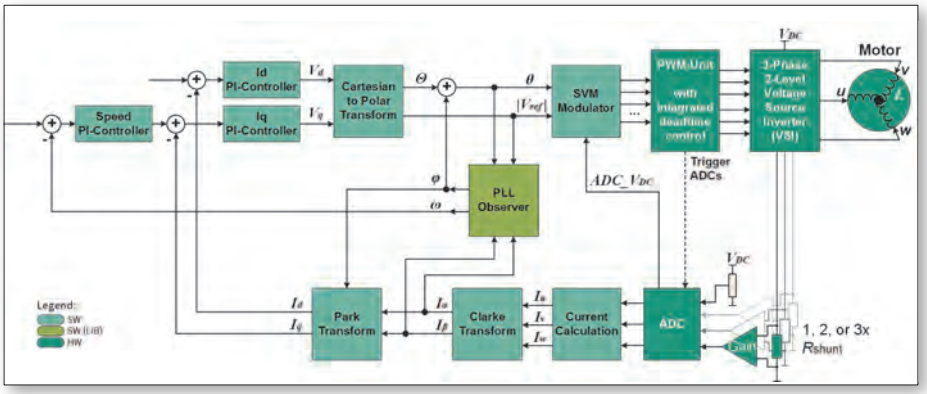

The XMC1302 supports efficient field-oriented control.

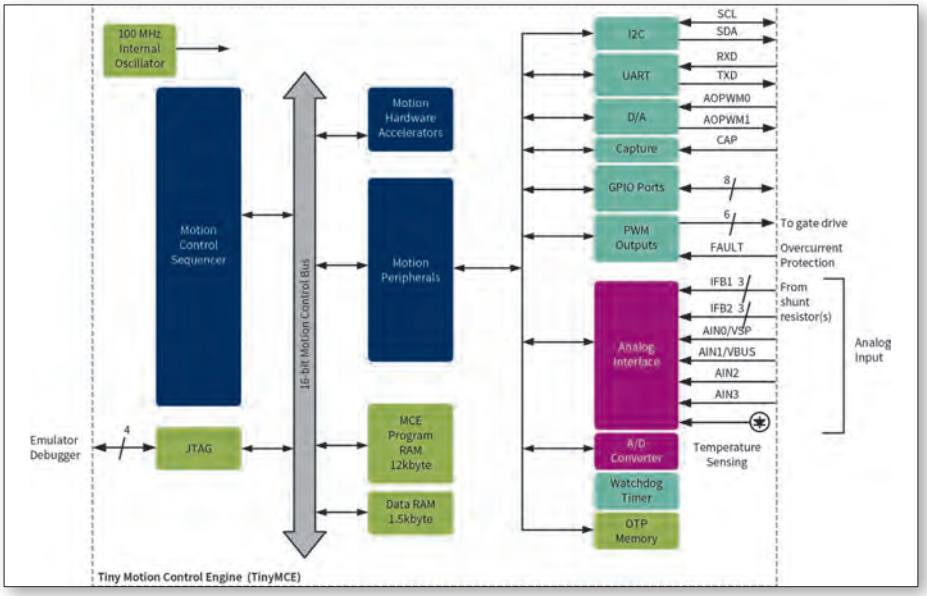

Alternative motor control with the IRMCK099M controller board, where a motor control ASIC with hardware-based, sensorless FOC control is implemented

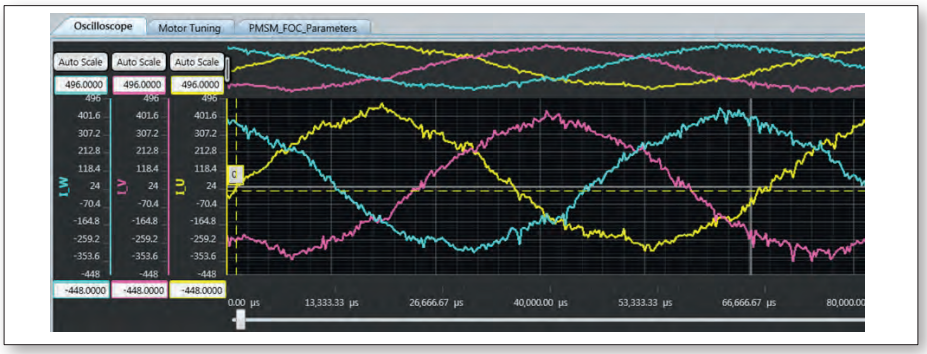

Using the µC/Probe XMC, it is possible to display critical control loops in the motor control in real-time with the aid of an 8-channel digital storage oscilloscope.

It computes in parallel to the main CPU, which in turn allows powerful algorithms to be implemented more easily and precisely for the field-oriented motor control. Another motor control specific block is the programmable POSIF interface, which enables using Hall sensors as well as incremental encoder for position sensing. Using the POSIF allows the accuracy to be improved and the software simplified for various applications, since the relevant data can be recorded simultaneously. A low-pass filter suppresses the noise and interfering impulses from the Hall sensor and rotary encoder, which could lead to incorrect position and/or speed readings.

The CCU8 supports applications with more complex PWM signal generation by means of complementary switches in half-bridge configuration and multi-phase control. These functions in conjunction with a highly versatile and programmable scheme for signal conditioning mean that the CCU8 is especially suited for high-performance motor controls as well as multi-phase and multi-level systems. Typical applications that benefit from this include 3-phase inverters for drives, 3-level inverters for solar modules and half-bridge converters. Using the CCU8 timer unit and the additional compare channel also enables the option of defining different dead times for rising and falling edges and generating asymmetric PWM signals.

The IRMCK099 is a dedicated motion-control IC designed specifically to address the complex requirements associated with advanced sensorless motor control applications. Unlike other motor control ICs that require some programming, the IRMCK099 features a Motor Control Engine (MCE) with control algorithms (based on standard library blocks) included in firmware along with a hardware accelerator. This MCE implements sensorless FOC for both interior and surface PMSMs using single or leg shunt current feedback through a combination of hardware and firmware elements. The IRMCK099 incorporates all major system elements in a 5 mm x 5 mm, 32-pin QFN package that operates from a single 3.3V supply. A built-in A/D converter offers 12-bit resolution and a 2µs conversion time, making it suited for precision applications. Alongside the advanced A/D converter is a 100 MHz internal oscillator that removes the need for an external clock. A Sigma Delta DAC provides a two-channel analog output and all analogue inputs are factory calibrated. Along with 16 kB of on-board OTP memory with a CRC memory check, the device features advanced communication and interface subsystems including a 57.6kbps UART, 8 digital GPIO ports, 6 PWM outputs and an I2C interface. Four register-selectable control inputs enhance design flexibility by supporting motor speed and direction control via UART, analog voltage, frequency or duty cycle. This ensures simplified interfaces within a wide variety of applications. The single shunt reconstruction feature allows the accurate current measurement needed for FOC control to be achieved with a single external shunt, thus minimizing external analog and digital circuitry. Phase shift PWM eliminates the minimum pulse limitation, improving motor start at low speeds and reducing acoustic noise in operation. Built-in safety and protection features include rotor lock protection and catch spin to detect any rotation of the motor before a control input is applied.

As opposed to MCU or DSP-based FOC designs, the dedicated iMOTION IC approach makes it possible to realize a motion control design without any programming. Control parameter calculations are performed using a GUI-based drive configuration and design

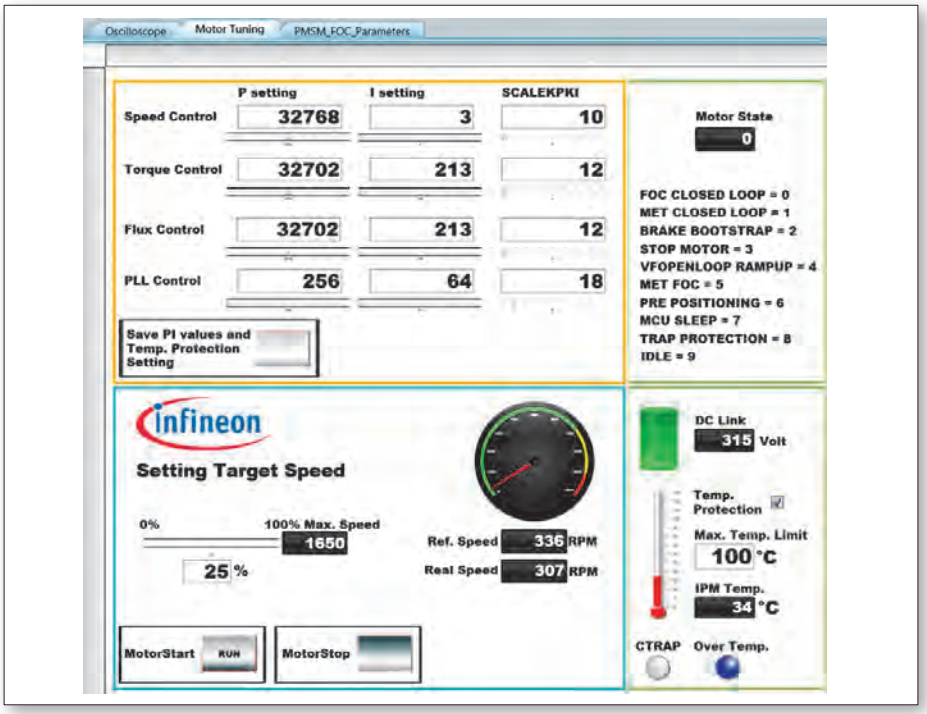

GUI-based software permits simple parameterization and tuning.

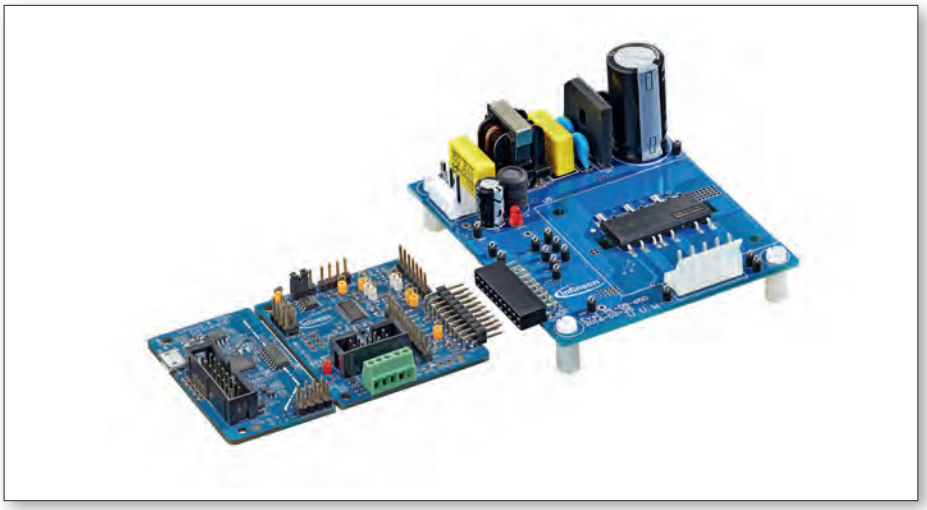

Complete design kits with two controller boards and four different power boards (250V and 500V)

tool (MCE Wizzard). The tool features a simple form-based dialog GUI that captures all the motor parameters and application information such as speed and acceleration in an easy-to-understand, engineer-friendly format – also explaining where to find the necessary information. Once calculated, the parameters are automatically exported to the MCE Designer tool. This tool facilitates motor drive control and testing as well as parameter tuning. A powerful parametric trace tool allows the user to trace and plot internal control variables, thus rapidly debugging and improving the motor drive performance. The iMOTION Modular Application Development Kit (MADK) is a compact and versatile evaluation system and a scalable development platform for 3-phase motor drives (115/230 V) currently covering the power range between 20W and 300W. It comprises controller and power boards for sensorless or optionally for sensor-assisted control. The kit can be used to implement a fully functioning motor system within less than an hour. All that is left for developers to do is to connect the boards to the PC, motor and power supply, download, install and parameterise the software – and the motor will already be running.

The iMOTION MADK kits offered address various motor control designs. Each kit comprises a control board with built-in or separate debug interface in addition to a complete power stage with a rectifier and EMI filter. Also included: motor control software (pre-installed or available for downloading) in addition to simple, GUI-based software (µC/Probe XMC or MCE Wizzard/MCE Designer) for the parameterization and tuning.

The µC/Probe XMC is a derivative of the µC/Probe from Micrium specially designed for XMC microcontrollers. This Windows-based application software allows the XMC memory to be written and read during operation without disruptions in order to optimise the application. With µC/Probe XMC users can generate their own GUI interface. It is possible to position and view the graphic components easily using drag and drop functionality. The global variables and memory contents can be monitored in real-time, also with the aid of a 4-channel digital storage oscilloscope. The oscilloscope can be simply enabled with an object code, which also determines the configuration of the sampling rate or the size of the buffer storage.

For the various applications or power ranges, it is currently possible to combine two different controller boards with six different power boards. One of the controller boards is based on XMC1302 MCU while a further control option is available with the IRMCK099M-based controller board. The power stages optimised for motor drives are based on the compact integrated power modules and available in 250V and 500V variants respectively. The modules comprise 3-phase inverters including high-voltage gate drivers and MOSFETs with low on-resistances.

If the XMC1302 controller card is used, the FOC software is either already installed or it can be downloaded via the Internet. The control card hardware is designed for also supporting either Hall sensors or innovative 3D magnetic sensors. The XMC1302 control card features on-board debugger with Segger J-Link technology. Further application software can be implemented via the free DAVE Integrated Development Environment (IDE) from Infineon, or other ARM IDEs, such as from Keil, IAR or Atollic.

The IRMCK099 control card is supplied along with a separate isolated debug interface card. The controller is delivered pre-programmed with sensorless FOC firmware. Both MADK controller boards and six different power boards can be purchased separately through Infineon or via distributors worldwide. There are also four different kits offered combining an XMC1302 control board with one of the four power boards of the µIPM IRSM836 or µIPM-DIP IRSM505 series (500V or 250V option). Infineon will be gradually extending its offering of controller and inverter boards, e.g. with up to 1kW. ■

Box PC with SMARC 2.0 module based on the Intel Apollo Lake processors

By Zeljko Loncaric, congatec

This article introduces an all-new Box PC developed as an IoT gateway. Not only is the system standard eNUC for Box PCs new, the SMARC 2.0 module standard is new and the Intel Atom, Celeron and Pentium processors that are being deployed in the system were just launched in October 2016.

Figure 1. The TeNUC-100 Box PC from Technagon was developed as an IoT gateway according to the eNUC standard and is based on SMARC 2.0 modules from congatec with brand-new Intel Atom, Celeron and Pentium processors (code name Apollo Lake)

■ Due to the increasing need for powerful solutions with IoT connectivity, Technagon made the strategic decision to develop its own IoT gateway to be used as an application-ready component in numerous customer applications. The inspiration for this development came from one customer, an automotive manufacturer, who had specific requirements for its charging stations at dealerships. A central system had to integrate several e-charging stations via a charge point server into the respective dealer network. The charge point server provides the dealer with access to all relevant on-site functions and allows the operator to carry out comprehensive remote management functions necessary for smooth operation of the charging stations, including functional upgrades. At the same time, the charge point server ensures a secure separation of the dealer network from the charging infrastructure.

In the past, Technagon developed and manufactured this type of gateway as 19-inch rack systems for server rooms of the dealers. Owing to the increasingly high performance of smaller embedded processors - such as the recently launched Intel Atom, Celeron and Pentium processors (code name Apollo Lake) - and the need for embedded systems in the outdoor area for individual charging stations or in digital signage products, the company decided to develop a rugged embedded Box PC family that can be deployed anywhere, i.e. in control cabinets, wall-mounted applications, outdoors (TeNUC-100 R) and even as a system on a Vesa mount installed at the back of a display.

The challenge was to cater for extremely diverse potential application areas while at the same time incorporating standardized basic technology. The mission was to develop a solution which is as innovative, forward-looking and comprehensive as possible and based on existing standards and therefore is able to reap the benefits of a comprehensive ecosystem of existing solution modules. As is repeatedly the case, the commercial sector delivered the blueprint for the embedded computer standard which best fits the requirements: embedded NUC, eNUC for short. The manufacturer-independent board and system standard developed by the SGET Standardization Group for Embedded Technologies e.V. specifies a 10.16 x 10.16 cm² board with a primary I/O area on the front and an optional one at the back. The cooling solution and power supply design are also specified to guarantee interchangeability of boards and housings. With a base area of just 100 cm², the board is predestined for small system designs. In future, the eNUC specification will also standardize the design of the housing so that users will be able to rely on a broad portfolio of system specifications, which will include system housing from Technagon. So, on the new TeNUC-100 system, everything is based on standards.

However, to qualify as an IoT gateway, the eNUC Box PC has to offer all the right interfaces. Along with two standard Ethernet interfaces, via which horizontal or vertical integration of the system into two separate nets is possible, wireless interface support is crucial too, as IoT applications are often connected via LTE or 3G/4G mobile radio interfaces to central cloud servers. Plus, they are also deployed as gateways for a whole range of different wireless sensor networks.

Radio standards for especially energy-saving long-range communication over several kilometres, such as LoRa, 6LoWPAN and Sigfox or 3GPP, LTE-MTC and UNB, are currently in great demand in smart city and smart energy nets. Additionally, for close-range communication wireless protocols, such as WLAN, BlueTooth (BTLE), NFC and other IoT radio protocols such as Zigbee, Z Wave or Thread as well as proprietary radio protocols are being deployed which also have to support a universally applicable IoT gateway. For this reason, the TeNUC Box PC is equipped with two flexible Mini-PCIe slots, which accommodate extension modules for all the mentioned standards. The system can be designed

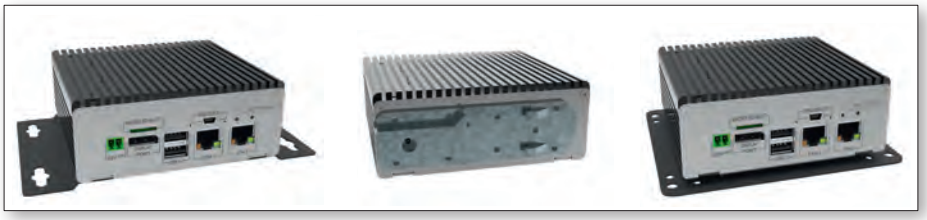

Figure 2. The TeNUC-100 Box PC can be installed flexibly as a desktop, a DIN rail PC or a wall and monitor mount and is also available as a TeNUC-100 R in a rugged version in the aluminium housing and with IP54 protection

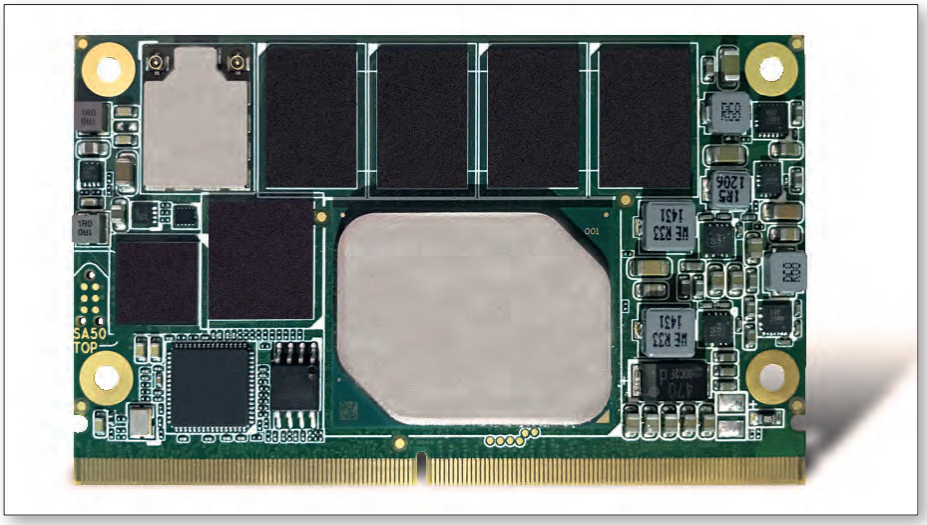

Figure 3. The congatec SMARC 2.0 module is equipped with Intel Atom, Celeron and Pentium processors (code name Apollo Lake) and suitable for outdoor use in smart city applications

with two antennas for each module to ensure highest radio quality for every wireless interface. As an option, a third radio interface (also equipped with two antennas) can also be integrated on the eNUC carrier board. The system can take advantage of the SMARC 2.0 standard-compliant innovative modules as this standard allows direct implementation of native radio interfaces on the module.

The choice of SMARC 2.0 is also technically the most powerful module concept, as it integrates both ARM and x86 technology, which offers decisive scalability and caters for variants to match a wide range of device requirements and developer preferences. With its credit card-sized design, it is the module standard with the most innovative and largest number of interfaces. SMARC 2.0 offers up to four display interfaces, 2x Gigabit Ethernet, PCIe, 2x MIPI-CSI camera inputs, 2x USB 3.0, 6x USB 2.0, 4x COM, CAN, SPI and I2C as well as HDA and 2x I2S for audio, facilitating the design of highly innovative IoT gateways in the Box PC format.

Technagon developed the eNUC standard Box PC based on SMARC 2.0, so the housing and carrier boards were developed in-house. For the required modules used in TeNUC-100 Box PC design, Technagon chose the embedded supplier congatec. As a sales technology partner for customer-specific system design and manufacturing services in the field of POS and ticket systems, vending systems, digital signage and eMobility, congatec technology became an integral part of the Technagon solution. The development service provider performance bandwidth ranges from application-specific carrier boards and embedded Box PCs, it can cover system integration of all components including the design and manufacturing of the customer-specific housing which can even culminate in the complexity and size of an e-mobility charging station or kiosk systems.

This close sales technology partnership with congatec enables Technagon to integrate latest processor technology into customer-specific projects in a very fast way, giving their customers an enormous head start in terms of technology and time-to-market. The Box PC is equipped with first congatec SMARC 2.0 Computer-on-Module with Intel Atom, Celeron and Pentium processors (code name Apollo Lake). These processors were only released at the end of October 2016 parallel to the launch of the first congatec SMARC 2.0 module. And parallel to this, Technagon developed the TeNUC-100 Box PC and was in a position to present the system coinciding with the processor launch. The feature set of the TeNUC-100 Box PC with a wide-range power supply is convincing. The basic configuration includes a Displayport, LVDS and MIPI CSI camera support for smart digital signage, vision-based access control, general video surveillance and other interactive video applications. IoT wireless interfaces can be provided via two Mini PCIe slots, which, of course, can also accommodate other extension modules. Connections via 2x LAN with Power over Ethernet help to reduce cabling costs as the system does not need its own power supply. 2x USB 3.0 are available for peripherals and a USB client port as a local management interface.

In terms of storage media, the following is available: a MicroSD card slot and 1x SATA - including power - as well as up to 64 GB flash memory on the SMARC 2.0 module. As an option GPIOs, serial RS232 and RS485 interfaces as well as I2S and HDA, 2x CAN, SPI, eSPI and security chips as well as further sensors for temperature, acceleration, rotation, etc can be integrated via an extension slot. Various designs of the extension slot board can also be used to create customer-specific variants. This feature set is available in a wide range of housing configurations, i.e. a box version for desktop deployment, two for wall or DIN rail mounting, as well as a variant with Vesa mount for monitor mount. Further to this, different housing designs are available, even including a robust aluminium version for outdoor applications, i.e. in smart cities, the energy industry, in the field of bus and rail as well as in a wide range of other IoT applications.

As Technagon often delivers to its solution provider clients, such as the already mentioned automotive manufacturer, the complete middleware right up to the application layer or even the entire application for charge point servers as well as a wide range of other systems, the TeNUC-100 Box PC does not come as a standard product. It is provided to ODMS customers and verified both on the software side and with the corresponding extension modules that the OEM customer requires for each particular Original Design and Manufacturing project.

The boxes offer a myriad of different possibilities for the customers. They can be provided with, for example, an IoT gateway for sensor networks in smart cities or for all types of vending machines, which today all require IoT connections to integrate new payment systems such as NFC or, for example, to even comply with fiscal requirements for accounting as efficiently as possible.

To fulfil these demands, in smart city projects the company integrates, for example, the distributed sensors and in the case of machines the complete peripherals and offers OEMs an

individually tailored entire solution including the middleware and applications layers up to the device cloud. Its complete Extended Engineering Service is therefore always a major aspect and fits the needs of many new IoT start-ups, who need to get their innovative, genuine and high-quality solutions to market as quickly and efficiently as possible. OEMs who purely need standard platform configurations without additional ODM services do not have to feel left out, as congatec decided to extend its IoT gateway offer to include complete solution platforms. The reason for this is that the large and rapidly growing IoT market is best served with the classic range of embedded computer boards and modules from congatec, but OEMs increasingly require IoT gateways as completely application-ready platforms. In order to meet this growing demand, congatec is working on a portfolio of different configurable gateway platforms, which can be optimized to meet a wide range of very specific application requirements. Together with the matching Embedded Design and Manufacturing Services from Technagon, the result is a comprehensive range of solutions which meets every specific OEM requirement for IoT gateways. ■

Product News

Hall-Stand 4858

TQ extends its x86 product portfolio

The COM Express Basic (Type 6) module TQMx70EB supports Intel Core i3, i5 and i7 7000E series processors (with up to quad core 3.7 GHz / 8 MB cache). For server-like applications the module supports also quad core Intel Xeon E3-1500 v6 processors. This new high-end module offers increased CPU and graphics performance. Equipped with 10bit codecs (HEVC/VP9) the module delivers high color fidelity premium graphics content. Latest SSD mass storage devices like Intel Optane with 3D XPoint technology are supported which results in less read/write latency and high data throughput. The TQMx70EB sets a new level of performance and addresses applications in gaming, medical, automation, simulation and data analyzing.

News ID 4858

Hall-Stand 4887

TQ: robustness throughout the embedded product range

TQ's appearance at embedded world 2017 spotlights the robustness of TQ embedded products. TQ modules cover not only standard solutions but also special applications for harsh industrial environments. A well established and thought out industrial design ensures that the modules function reliably throughout their whole life-cycle. TQ's in-house production guarantees optimal manufacturing quality. TQ modules can even meet increased requirements, such as condensation or other aggressive environmental parameters, thanks to their conformal coating.

News ID 4887

Hall-Stand 4959

EKF: quad M.2 WWAN module carrier

EKF introduces the SPX-PHASE, a peripheral slot card for CompactPCI Serial systems. The board serves as a quad carrier for M.2 3042 style WWAN modules, e.g. 3G or 4G (LTE) modems, and also forthcoming technologies beyond 4G. Four front panel Mini SIM card slots are provided. The R/F module connectors are strapped to an on-board 4:1 R/F combiner circuit, for minimum external antenna cabling effort. The SPX-PHASE is well suited especially for higher throughput applications which require WWAN bandwidth aggregation and redundancy.

News ID 4959

Hall-Stand 4950

congatec: latest OS implementation further simplifies development of IoT-connected devices

congatec has announced its comprehensive support of the IoT editions of Windows 10. These full versions of Windows 10 are characterized by their extended security functions and lockdown options for IoT-connected appliances. congatec has integrated these innovative features on all its current boards with Intel Atom, Celeron, Pentium, Core and Xeon processors as well as AMD Embedded R and G Series processors. Customers benefit from an overall simplified development of IoT connected devices.

News ID 4950

Hall-Stand 4951

PEAK-System: CAN FD interface for M.2 slot

To this year's embedded world 2017, PEAK-System present a CAN FD interface for the M.2 slot. The plug-in card called PCAN-M.2 will be available as a single, dual, and four-channel version. Furthermore, the PCAN-miniPCIe FD, the four-channel PCAN-PCI Express FD, and the opto-decoupled PCAN-USB X6 with M12 circular connectors and the ingress protection IP64 are presented for the first time. Thanks to this increased protection and the six CAN and CAN FD channels, the M12 USB adapter is especially suitable for motor vehicle test stands and HIL simulations harsh environments.

News ID 4951

Apollo Lake processors will be the heart of future embedded computers

By Peter Ahne, Portwell

The embedded industry must wait for the latest Intel process generation - code name Apollo Lake – but has developed numerous boards and modules based on this CPU series. Examples of different standards and form factors from COM express to 3.5 inch document the versatility of the building block and the high interest of the market.

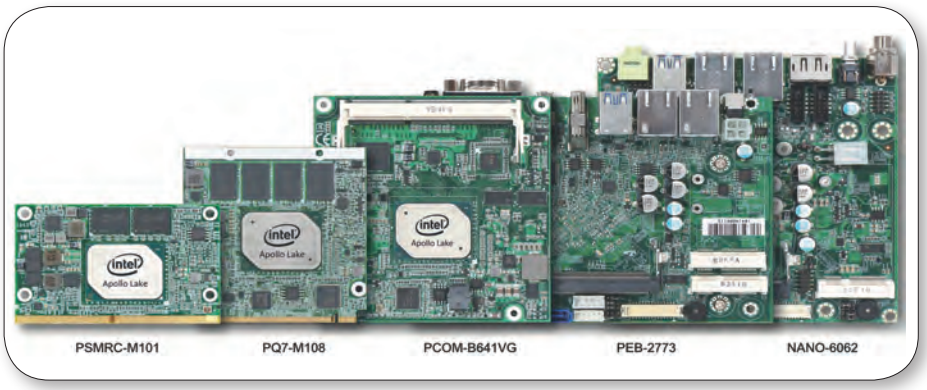

Based on the five Apollo Lake processors on the embedded roadmap, Portwell has already launched computer-on-modules and boards with these processors.

■ Almost 13 years ago Intel launched the Pentium M, and made the transition from the pure hunt for more performance by constantly raising the clock frequency, as in the days of the Pentium x processors, to the age of less power consumption. In 2006 the Tick-Tock model was announced, with its continuous improvement of production technology with a reduced nm process – Tick and Tock standing for new micro architecture. At that time some may have asked themselves, when will Intel hit the physical limits? The answer is in ten years. With the fourth generation of the Core i processors it has happened, and instead of a Tick, a so-called refresh is established. Refresh stands for everything (production technology and micro architecture) staying the same, but the processor getting improved internally, making performance increase possible. Intel's new development process is called PAO: Process – Architecture – Optimization.

■ For the Atom processor series, the old Tick–Tock model is still in place. The last step, a Tock-generated Apollo Lake, is the latest in the successful Atom story. Time will show when the PAO development model will reach Atom processors as well. But for the moment, the market is watching what boards and modules based on Apollo Lake processors will come. The tapered edges of the metal cap and the metal cap itself are a hallmark for the Intel Apollo Lake processor – but in everyday life it will disappear under the heat sink. Offering a broader contact space for cooling solutions is one feature that will help embedded vendors to design more compact systems with even more processing power. This is the superficial improvement, but what improvements have been made inside the processor? The newcomers are manufactured in 14nm technology and provide approximately 30 percent better computing and 45 percent more graphics/video performance than their predecessors. The power consumption remains relatively low at 6 to 12W TDP and stays at the level of the former generation.

The new silicon features Intel Gen 9 graphics, which in the past was reserved for the sixth-generation Core processors only. This gives the new CPU series a significantly improved support for 4k graphics resolution. Up to three independent displays can be operated simultaneously, and DirectX 1.2, OpenCL 2.0 and OpenGL 4.2 are supported. The integrated security engine provides better security while the optional Error Correcting Code (ECC) preserves data integrity. The protection capabilities in the new processors have been built from the ground up to ensure a new level of security. Each time the system starts up secure boot with Intel Platform Trust Technology helps to keep the device safe, blocking dangerous programs, so only trusted software is launched. All this works behind the scenes of the new processor but delivers more power and advanced cooling to applications.

There is just one weakness with the new Apollo Lakes: not all five types on the embedded roadmap are available in quantities right now. The embedded world must still wait for the Atom x5 and Atom x7 versions. For this reason, most versions of the already designed boards and modules are delayed until Intel starts to supply the new Atom version. Based on these five processors on the embedded roadmap, embedded vendors have already launched computer-on-modules and boards based on these processors. As in the past, the Atom processors are mainly designed on small form factors or one with limited power envelope. Two Apollo Lake based modules and two boards will be taken as example to showcase the variety of products the Apollo Lake potentially offers.

The ultra-low-power Intel Atom processor E3900 family as the engine of a COMe basic module allows to make a design that brings COM units with a power consumption under 6W~12W which is a perfect fit for fanless applications. This also allows supporting a wide -40°C to 85°C industrial grade temperature range. On the memory

Product Name	Atom™ x5-E3930	Atom™ x5-E3940	Atom™ x7-E3950	Celeron® N3350	Pentium® N4200
Cores/Threads	2	4	4	2	4
Base/Burst Frequency	1.3/1.8 GHz	1.6/1.8 GHz	1.6/2.0 GHz	1.1/2.4 GHz	1.1/2.5 GHz
TDP	6.5 W	9.5 W	12 W	6 W	6 W
Max Memory	8 GB	8 GB	8 GB	8 GB	8 GB
Memory Types	DDR3L up to 1866 MT/s; LPDDR4 up to 2133 MT/s	DDR3L up to 1866 MT/s; LPDDR4 up to 2133 MT/s	DDR3L up to 1866MT/s; LPDDR4 up to 2400 MT/s	DDR3L/LPDDR3 up to 1866 MT/s; LPDDR4 up to 2400 MT/s	DDR3L/LPDDR3 up to 1866 MT/s; LPDDR4 up to 2400 MT/s
Max Memory Channels	4	4	4	2	2
ECC Supported	Yes	Yes	Yes	No	No
Graphics ‡	Intel® HD Graphics 500	Intel® HD Graphics 500	Intel® HD Graphics 505	Intel® HD Graphics 500	Intel® HD Graphics 505

Apollo Lake processors on the Intel Embedded Roadmap

side 32GB DDR3L 1600/1866 MT/s SDRAM on two 204-pin SODIMM sockets makes it faster than Baytrail or Braswell-based modules. Expansion interfaces support four PCI Express x1 Gen2 (5.0GT/s) (maximum) for enhanced video performance and offer the flexibility to configure to two x2 lanes or one x4 lane. In addition, three high-resolution display interfaces are supported: DP (DisplayPort) or HDMI, VGA and LVDS with double 3D performance compared to the previous generation. The whole package provides a more powerful solution for targeting IoT and Industry 4.0 applications.

Designed according to the latest improved and particularly rich feature set on SGeT SMARC 2.0 specification, SMARC 2.0-based modules are capable of offering protection against unauthorized re-use or re-provisioning of original parts. Modules like this can be equipped with 2x USB 3.0, 6x USB 2.0 (1x OTG), 1x SDIO3.0, 1x SATA 3, and up to 4x PCIe Gen2 lanes. Also onboard are 5x I2C, 2x SPI, 4x UART, 12x GPIO, and HD audio featuring an integrated audio DSP with Intel Smart Sound Technology. The modules can handle industrial -40 to +85°C temperatures.

The NANO-ITX form factor is quite famous in the embedded market because of its great balance of small size and reasonable connectivity set. Equipped with the Apollo Lake such boards are able to operate with a thermal design power (TDP) under 12W for fanless applications and also support a wide industrial temperature range from -40°C to 85°C and a wide voltage range of power input from 12V to 24V for rugged applications. The flat/low-profile design - measuring 16.4mm in height with I/O shield - allows space-saving installation in display and panel PCs, making the realization of digital signage and control solutions for industry and business applications a quick and easy task. In addition, they support up to 4096x2160 pixels with DisplayPort, 1920x1200 pixels with VGA, and dual channel LVDS up to 1920x1200 pixels, and are designed with the flexibility for connecting to three independent display interfaces. Support for four USB 3.0 ports ensures fast data transmission with low-power consumption. One 5 Gb/s PCI Express 2.0 lane can be used as 1x full-size mini-PCIe. Two SATA 3.0 interfaces with up to 6 Gb/s (one of them available as mSATA and the other for SATA) allow quick and flexible system expansions. The Intel I210-IT Gigabit Ethernet controllers provide dual Gigabit Ethernet LAN access via the two RJ45 ports. A significant benefit of using NANO-ITX is that customers can design their own unique systems for medical, networking, Panel PC, kiosk and digital signage applications based on it.

A 3.5' SBC equipped with the Apollo Lake processor really takes advantage from its processor power and feature set. On the graphic side, it supports one dual-channel 24-bit LVDS connection, one DisplayPort (DP) and HDMI on rear I/O with resolution up to 4096x2160. The 204-pin non-ECC SODIMM provides maximum memory, making it capable of supporting up to 8GB of DDR3L. It can support a total of six COM ports, six USB 3.0 ports and dual Gigabit Ethernet. The thermal design power (TDP) under 12W allows fanless

applications. The board also supports a wide industrial temperature range from -40°C to 85°C and a wide voltage range of power inputs from 12V to 24V for rugged applications. Two 5Gbit/s PCI Express 2.0 lanes can be for use as 1x full-size Mini-PCIe (supporting mSATA) and 1x half-size Mini-PCIe socket. It also supports one SATA 3.0 interface with up to 6 Gb/s to allow quick and flexible system expansions. The Intel Ethernet Controller I210-IT provides dual Gigabit Ethernet LAN access via the two RJ45 ports. Such a board fits into a variety of market segments and applications such as retail, networking, Panel PC, industrial automation and digital security surveillance. The embedded community will continue to extend the portfolio on board level form factors equipped with the Apollo Lake processor family. And soon we'll see the first Box PCs and Panel PCs based on Apollo Lake as well. ■

Product News

Hall-Stand 4905

DFI: rugged small form factor Embedded computer

DFI Tech offers a new small form factor rugged system to satisfy mission-critical and cost-efficiency needs. The EC500 from DFI Tech features a 6th/4th generation Intel Core i7/i5/i3 processor. The conduction-cooled system supports dual 8GB DDR3 memory, a 2.5" SATA drive bay with SATA III capability, and 1 mSATA module via the Mini PCIe socket. The unit also facilitates rich graphics capability with 24-bit color depth up to 1920x1200 resolution and comes with Intel HD Graphics.

News ID 4905

Hall-Stand 4940

Concurrent: support for latest gen Intel Xeon processor E3-1500 v6

Concurrent Technologies announces support for the latest generation of Intel Xeon processor. The first product from Concurrent Technologies based on this new device will be a new AdvancedMC module, AM G6x/msd, which is suitable for high-speed physics experiments, instrumentation and test based applications. AM G6x/msd customers will be able to specify a choice of processors including the Intel Xeon processor E3-1505M v6 or Intel Xeon processor E3-1505L v6.

News ID 4940

Hall-Stand 4908

u-blox: ultra-small multi-GNSS module

u-blox announces the launch of ZOE-M8G an ultra-compact GNSS receiver module, especially designed for markets where small size, minimal weight and high location precision are essential. ZOE-M8G offers exceptionally high location accuracy by concurrently connecting to GPS, Galileo and either GLONASS or BeiDou. It also provides 167 dBm navigation sensitivity. This makes the ultra-small ZOE-M8G perfect for wearable devices, unmanned aerial vehicles and asset tracker applications.

News ID 4908

Video Surveillance

with High Capacity Storage & Max I/O


City Surveillance


Banking and Finance


Critical Infrastructure


Hospitals


Retail & Commercial


Transportation


X10DRX Motherboard
11x PCI-E Slots (10 x8 & 1 x4)


SC846XE2C
4U 24x 3.5" Drive Bays w/BBP


SYS-6048R-TXR
4U 5x 3.5" Drive Bays


SYS-6038-TXR
3U 8x 3.5" Drive Bays


SYS-2028R-TXR
2U 16x 2.5" Drive Bays


SSG-6048R-E1CR60N
4U 60x3.5" Drive Bays


SSG-6038R-E1CR16N
3U 16x3.5" Drive Bays


SSG-6028R-E1CR24N
2U 24x3.5" Drive Bays


SYS-5029S-TN2
Mini Tower 4x3.5" Drive Bays


SYS-E200-9B
Compact Box System

Low Power Intel® Quark™, Intel® Atom™, Intel® Core™ processor families, and High Performance Intel® Xeon® processors

- Performance Data Analytics
- Security Fast Data Encryption
- Media Transcoding
- Storage Expansion
- **Video Surveillance with Max I/O**
- Remote Management OOB
- Digital Video Solutions
- Analog Video Solutions

Intel Inside®.
Powerful Productivity
Outside.

TITANIUM

We Power the Cloud™

Learn more at www.supermicro.com/embedded

CompactPCI Serial now reaches out into space

By Manfred Schmitz, MEN

Space CompactPCI Serial allows the use of established industrial technology in space. This permits the cost-effective use of the latest technology for space applications. The open standard guarantees the interoperability of different boards from different suppliers, and helps re-use solutions from mission to mission.

Figure 1. A typical application for Space CompactPCI Serial is for example the implementation of the platform and the payload controller on board satellites.

■ Recently, the US-based Internet service supplier OneWeb ordered 900 satellites to provide global Internet broadband service. Knowing that this volume is more than half of the total 1400 satellites already in orbit, and the cost of sending one up into space is about 100 million dollars, you need to start thinking about new computer technologies that could help to manage this mass of satellites that need to be produced every year.

The commercialization of aerospace, like this OneWeb project, progresses and increases the pressure on established companies to develop competitive products for this market. One possible option is to count on standards in order to save time and cost during development, and also during project life-cycle. Taking an existing standard and adapting it to the requirements for space was also the solution a group of global players in space like Airbus, Thales Alenia, and others went for. They decided to go with a rather new and powerful, already industry-proven with its predecessor versions, standard: CompactPCI Serial. Using a standard like this also helps to contribute to the obsolescence problem of avionics hardware, as well as to significantly reduce complexity and costs for hardware/software. The new technical PICMG sub-committee Space CompactPCI Serial is now extending the current CompactPCI Serial specification by a sub-standard covering the specific requirements for space applications. Players in this working group are, in addition to Airbus and the initiating companies STI Spacetech, Sysgo and TTTech, also Amphenol FCI, EKF, Elma Electronic, Erni, Fastwel, Harting, Heitec, Intel, Keysight, Pentair and Positronic, with MEN Micro acting in a consultative role, having already been strongly involved in the development of the CompactPCI PlusIO and CompactPCI Serial standard.

Even if VPX was – with the exception of completely customized solutions – for a long time the only standard for embedded systems in space, CompactPCI also has a long history in space applications. Famous examples are the Curiosity Mars rover for satellite control, or implementations for scientific tasks in the ISS. Even the CompactPCI Serial base specification has the mechanical and conduction cooling technologies needed for space already defined and in place.

In addition, compared to VPX, the flexible CompactPCI Serial standard offers an easier and cost-optimized development. Space CompactPCI Serial is therefore the most logical choice for implementing high-tech solutions for a highly sophisticated market, while both re-using and evolving proven industrial technologies and reaching significant cost reductions. At the same time, some unused features have been removed to make the standard more streamlined, and other properties have been added to optimize the standard for use cases. The two main changes in the extension of the CompactPCI Serial specification are the definition of a dual-star architecture for increased availability, and allowing the integration of different communication protocols common in space applications for both the dual-star and the full-mesh network (which was formerly restricted to Ethernet only).

The CompactPCI Serial basic specification defines a single-star architecture, while Space CompactPCI Serial now doubles these interconnects in a symmetrical way, so if one CPU fails, the functionality of the complete system will not be affected. Having this increase in availability was essential for use in space, especially since you cannot simply exchange a CPU card installed in a satellite in orbit. In addition to the system slot (A) on the left side of the system, a second system slot (B) on the right side of the system uses the same routing method. All seven peripheral slots are connected to both system slots, and both system slots are connected to each other. All together, these links build a fully meshed interconnectivity network. The full-mesh network is not restricted to any particular protocol, and may be used for physical interconnection standards like Ethernet, SpaceWire, TTEthernet, EtherSpace etc. Parallel to the full-mesh network,

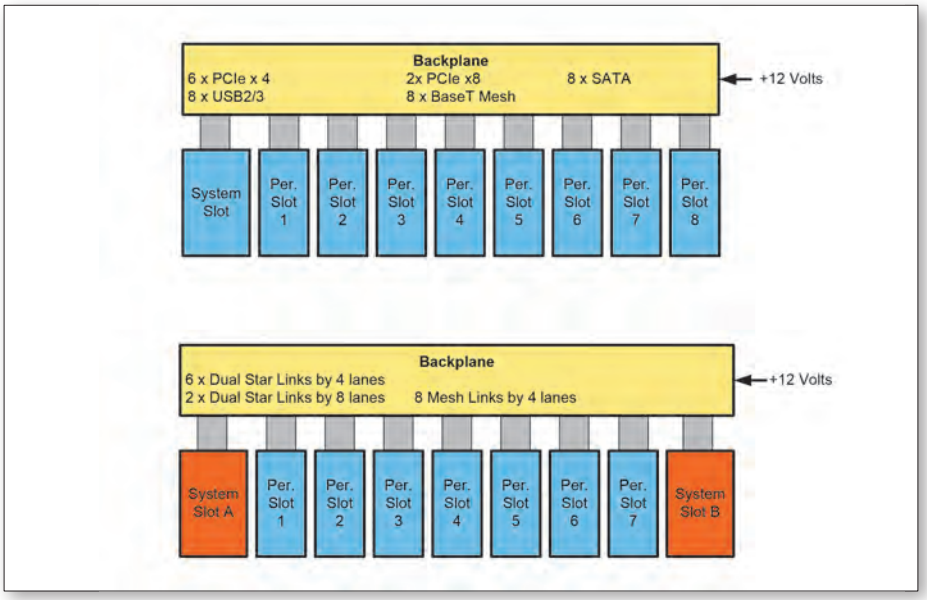

Figure 2. Space CompactPCI Serial defines a second system slot, forming a dual-star architecture for PCI-Express.

both system slots can be connected to any of the peripheral boards by means of eight specific differential links. These links could be also used for any protocol, depending on the application and the individual boards. This dual-star architecture is intended to be used in high availability solutions. The result is a parallel and flexible usage of the full-mesh Ethernet network via the backplane, as well as the dual-star architecture via PCI-Express or any other protocol.

In addition to that, the specification defines a utility connector, which can be controlled and configured via an open management bus. It takes over the hot-plug functionality, as it was used for CompactPCI Serial, and allows single cards within a system to be switched on and off. Hot-plug functionality is indeed not necessary for satellites in the orbit, but is extremely useful during integration on ground and for test systems, which can be still realized via PCI-Express and with common CompactPCI Serial cards.

As Space CompactPCI Serial is intended to be used in a conduction-cooled environment, the board-to-board pitch is defined to 5HP (= 25.4 mm) to allow a conduction cooling frame for each board. To make the standard specification easier, the pitch will remain the same even if the boards do not need a conduction cooling frame (e.g. for testing systems on ground), which would just require a 4HP pitch. At the moment, members of the working group are already working on specific standard backplanes for Space CompactPCI Serial. The high-speed backplane interconnects and the connectors are intended to support data rates of 12.5 Gbit/s per differential pair. The accelerated bandwidth of the full mesh is 400 Gbit/s, additionally the dual-star interconnect simultaneously supports a 1 Tbit/s throughput rate. Infrastructure signals allow comfortable and flexible system management. In addition to I^2C bus, CAN bus is also supported. The power supply is 12 volts only. Optionally, a 5 volt standby voltage could be used to support suspend modes or sleep modes. Power rails can be a single plane, or every board could be supplied and controlled individually. The power supply and the system management are not a part of the Space CompactPCI Serial specification.

The harsh environments in space and especially the vacuum conditions make high demands on the connectors. The material must not outgas, as some outgassed materials can deposit on the sensitive lenses in the satellites. A corresponding outgassing test has already confirmed the qualification of the connectors used for Space CompactPCI Serial. However, other mechanical or environmental measures, like SEU-resistance are not defined in the standard specification, as this always depends on the customer requirements, and the boards design within the application and end system. The working group just finished the specification recently and will plan to bring it into the PICMG ballot during the second quarter of 2017.

Space CompactPCI Serial allows the use of established industrial technology in space. This extends the cost-effective use of the latest technology to space applications. The open standard guarantees the interoperability of different boards from different suppliers, and helps re-use solutions from mission to mission. Keep an eye out for this new great standard, which open the doors for highly available and reliable applications, both in space and on the ground. ■

Facial recognition technology for next-generation fish-count surveys

This article is contributed by ADL

In the latest generation of the CAM-Trawl system developed by the Fish Management Acquisition (FMA) program, ADL worked with NMFS Alaska Fisheries Science Center making key improvements to the computer portion of the CAM-Trawl system, including higher compute performance, standards-based vision software, more rugged enclosure and user-friendly LED indicators.

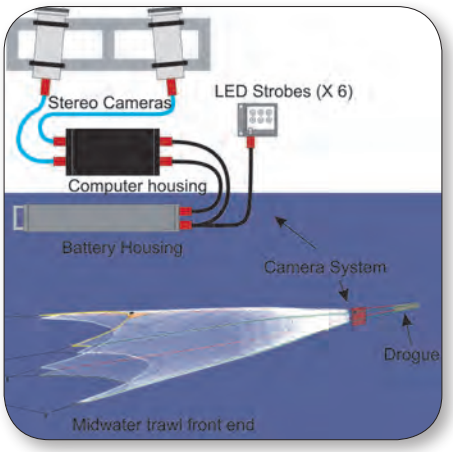

Figure 1. First generation CAM Trawl Picture

■ The conservation and management of fish stocks relies heavily on fish-count surveys. In an effort to improve the shortcomings of these surveys, NOAA (National Oceanic and Atmospheric Administration) and its National Marine Fisheries Service (NMFS) developed a camera-based trawl (CAM-Trawl) to bring the power of image capture and facial recognition technology to best address this problem. The Alaska Fisheries Science Center is the research branch of the National Oceanic and Atmospheric Administration National Marine Fisheries Service responsible for research on living marine resources in the coastal oceans off Alaska and parts of the West Coast of the United States. This region of nearly 3 million square miles includes the North Pacific Ocean and the East Bering Sea which support some of the most important commercial fisheries in the world. These waters are also home to the largest marine mammal populations in the nation.

The conservation and management of fish stocks relies heavily on fish-count surveys. Traditionally, this has been done through the use of fishing trawlers set up to do acoustic or bottom trawling which cast enormous nets into the sea to capture as many fish as possible over a specified area. These fish are then brought on-board and manually counted to develop the trawl survey. This methodology suffers from a number of shortcomings: errors associated with extrapolating from a small area survey to large area estimates, errors associated with extrapolating from the relatively long period of time necessary to gather and count fish over a region to real-time characteristics of the fish population, the relatively high cost of time and equipment for manual trawl surveys, the fact that some small fish species of interest are too small to be caught by the net mesh, and lastly, the fact that killing fish in the course of a trawl survey affects local fish populations which can sometimes be dangerously low.

As a result, NOAA (National Oceanic and Atmospheric Administration) and its National Marine Fisheries Service (NMFS) have been working for a number of years on a camera-based trawl (CAM-Trawl) technology to bring the power of image capture and facial recognition (loosely termed, Fishal Recognition) technology to best address the problem. Early generations of this CAM-Trawl system used the lengthy process of image capturing, storing images, and then removing the storage media for image analysis at some later time back at the lab.

In the latest generation of the CAM-Trawl system developed by the Fish Management Acquisition (FMA) Program, ADL Embedded Solutions worked closely with the NFMS Alaska Fisheries Science Center to make key improvements to the computer portion of the FMA CAM-Trawl system. These include the following. Upgrading to GeniCam-compliant vision platform to take advantage of the latest camera technology and image recognition algorithms. Upgrading to marine quality system enclosure with IP-67 rated ingress protection. Upgrading to a quad-core Intel Core i7 processor to enable real-time processing of image data. Design and building of a rugged small form factor system for mounting flexibility and ease of storage and transportation. The resulting FMA computer solution from ADL has helped achieve all the stated goals for this project. Namely, real-time fish-count data analysis capability for NOAA marine researchers to aid in their fish conservation efforts is now a reality.

The new system includes an ADL main control computer shown at left and an ADL image acquisition computer on the right in figure 2. System features include: stainless steel, water-proof (IP67) enclosures, Intel Core i7 processors, IP67-rated water-proof circular connectors with protective caps, LED activity lights for various functions, removable drive assemblies, and Topside Handrail Mounting Adapter. The ADL image acquisition computer (Vision Box) is responsible for: camera interface using up to 6x GigE ports, image capture,

Figure 2. The housed ADL system

real-time image recognition using processing algorithms on the Intel Quad Core i7 processor, and image storage including metadata for data/time, fish identity, etc.

The ADL control computer typically resides in the wheelhouse. It will steadily monitor a number of external sensors including geo data, time, pressure sensors, RFID tag readings and will orchestrate one or more Vision Boxes based on sensor input. Remote power on/off, clock syncing, start/stop image acquisition, etc is all dictated by the ADL control computer in a one-to-many relationship.

The results for the new FMA system have been very positive after experimental trawl surveys in 2016. Work is now underway to standardize the key components of the new system as a means of promoting the adoption of this real-time Fishal Recognition FMA system on many more surveys in the near future. ■

Product News

Hall-Stand 4864

Vecow: workstation-grade Intel Xeon/ Core i7 rugged vehicle computing system

Vecow launches the IVH-9200 Series Fanless Vehicle Computing System. Flexible LGA1151 Socket supports workstation-grade Intel Xeon/Core i7/i5/i3 processor (Skylake-S) running with advanced Intel C236 chipset and dual-channel DDR4 2133MHz ECC memory, up to 32GB capacity; Advanced Intel HD Graphics 530 supporting DirectX 12, OpenGL 4.4 and OpenCL 2.0 API, onboard DVI-I, DVI-D and DisplayPort display interfaces for Ultra HD 4K resolution. Vecow IVH-9200 makes new generation CPU performance, power efficiency, and graphics performance possible. PCIe 3.0 (8GT/s), Multiple SATA III (6Gbps), USB 3.0 (5Gbps), PoE (1Gbps) LAN and multiple wireless connections make seamless high-speed data conveying possible.

News ID 4864

Hall-Stand 4922

MEN: CompactPCI PlusIO SBC with Intel Atom Apollo Lake-I

MEN presents the latest member of its CompactPCI Intel CPU product family, which has been growing for more than eleven years now. The F26L is based on Intel's latest Apollo Lake-I platform with either two or four cores, and improved graphics options. As usual, the CompactPCI PlusIO board is backward compatible to all predecessors of this family. Due to the low-power architecture of its processor, the CPU board has a total power consumption of max. 6.5 to 12 watts, while having a clock frequency of up to 1.6 GHz.

News ID 4922

Hall-Stand 4901

Conrad adds Rittal's industrial grade compact enclosure AE products

Conrad Business Supplies has announced the immediate availability of Rittal's compact industrial grade AE series enclosures. The versatile, IP66 and NEMA 4 quality certified AE series is a premium product range offering a large mounting surface, easy fixing and secure cable entry. The enclosures are ideal applications in the electrical and industrial sectors.

News ID 4901

Hall-Stand 4869

Vector enables flashing of Linux-based ECUs according to OEM requirements

Vector is supplementing its portfolio of basic software products by adding a flash bootloader for Linux-based ECUs. Developers will now benefit from a quick and easy solution for reprogramming these ECUs according to OEM-specific requirements. It enables reliable updating of shell scripts and graphic data as well as configuration and operating system files.

News ID 4869

System design enablement and verification for the IoT

By Frank Schirrmeister, Cadence

This article introduces some new additions to the Cadence Verification Suite, which makes it the go-to solution for verification and software development for designs in the new age of the Internet of Things.

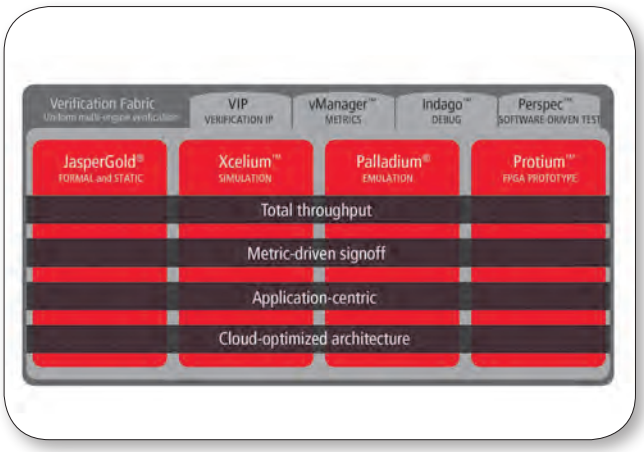

Figure 1. The Cadence Verification Suite

■ Spanning across most application domains, the internet of things (IoT) changes the requirements and priorities for hardware/software development, extending them well beyond the classic considerations of just performance, power and cost. Integrated verification - smartly combining formal verification, simulation at the transaction and classic register transfer level (RTL) - emulation and FPGA-based prototyping have become key requirements. The Cadence Verification Suite already led the way with best-in-class engines for formal verification and emulation. In February 2017, Cadence rolled out next-generation releases for simulation and FPGA-based prototyping. The Cadence Xcelium Parallel Simulator introduces the next era of simulation with true multicore parallelism, integrating technology from the 2016 Rocketick acquisition. The Protium S1 FPGA-based prototyping platform marks a new milestone by further reducing the time to prototype from months to days. It offers congruency with emulation and unique new features for software development, previously unknown to FPGA-based prototyping. With these new additions, the Cadence Verification Suite becomes the go-to solution for verification and software development for designs in the age of IoT. Designs labelled as IoT are applicable from wearable electronics connected to mobile devices, through smart homes connected to set-top boxes, the connected car including driver assistance and car-to-car connections, all the way to enterprise automation for smarter industrial, healthcare, city and energy management. They include designs for the edge node, like our fitness trackers and sensors in industrial automation, and designs for the hubs that aggregate data, like our cell phones, cars and home automation systems, as well as very complex designs for networking and servers in which data is transmitted and analyzed, enabling never-before-seen applications through machine learning and big-data analytics.

While in the past power, performance and cost were the main drivers helping design teams to prioritize designs, in the age of IoT they are now joined by other priorities like security, connectivity and in-field upgradeability. Not only are there more priorities to consider and properly balance, they also change depending on the end application. For fitness trackers, for instance, cost is highest priority, followed by power and connectivity. In contrast, for in-body health implants like defibrillators or cochlear implants for improved hearing, security and safety become the highest priority as nothing can go wrong in a device like this, ever. In-field upgradeability comes second—surgeries should not be repeated—followed by energy consumption. In reactive applications for automotive, industrial automation and robotics, performance comes first, followed by connectivity and in-field upgradeability. A lagging response in a car warning to the driver can easily put human life at stake or risk large revenue loss when industrial production may be brought to halt. Within the city infrastructure, for items like traffic lights or city-wide installations of important control functions, in-field upgradeability becomes the highest priority due to the volume and distribution in the field, followed by security of the infrastructure and cost.

Given the varying development needs for edge nodes, hubs, networks and servers, trying to meet varying priorities across multiple application domains, the vehicles for verification and software development need to be extremely flexible and require close interaction. Figure 1 shows the Cadence Verification Suite with its four core engines of formal, simulation, emulation and prototyping. It is crucial for development teams to be able to efficiently move a design through the different engines, and to capitalize on the individual strength of the engines by optimally combining them. A verification fabric spans across the engines and provides a unified user experience for key tasks like verification management, debug and portable testing. In combination, engines and fabric are geared to

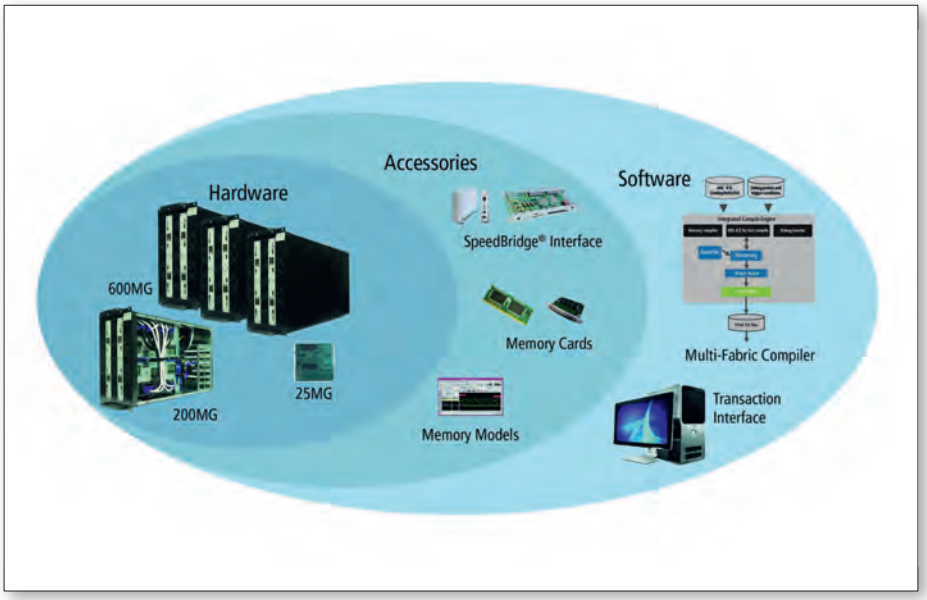

Figure 2. Cadence FPGA-based prototyping portfolio

optimize total throughput for designs, allowing fastest time-to-market and to allow metric-driven design, giving developers a clear indication when they are ready to proceed and roll out their products. The resulting software development and verification flows are optimized for specific application domains, such as adding specific capabilities around functional safety and ISO 26262 compliance for automotive, and are architected to be remotely accessible through cloud-based design.

Focusing in on the hardware engines, emulation and FPGA-based prototyping have long been available to accelerate speeds, extending the range of verification into hardware/software verification, software development and system validation. In the era of IoT, software development clearly has become the long pole in the tent, gating time to product delivery, and with that time, to revenue. Being able to develop software early and in parallel to hardware has become key to product success and requires the speed of hardware-assisted development.

Emulation, typically in the MHz range, is classically focused on high-value use models extending hardware verification to low-power optimization, performance analysis and even early verification of test sequences for post-silicon verification, all enabled by simulation-like debug, fast bring-up and very flexible allocation of tasks in emulation farms. With its higher speed in the tens of MHz or even 100MHz, the FPGA-based prototyping sweet spot has classically been software development and system validation. Neither engine has been without challenges. Emulation has always been somewhat speed-limited, and while smart connections to virtual platforms in hybrid setups can dramatically accelerate the time to point of interest by 50X, FPGA-based prototyping generally achieves the appropriate performance that satisfies the needs of software developers. However, time to prototype, due to the largely manual optimizations required and the need to re-write the RTL to make it FPGA friendly, has traditionally been long, often taking months.

The Protium S1 platform, together with the Cadence Palladium Z1 Enterprise Emulation Platform, directly addresses these challenges and revolutionizes bring-up time by an average of 80% from months to weeks or even days. It gives users the fastest path to prototype for software development, system validation and hardware regressions.

To enable this, the key innovations delivered in the Protium S1 platform are on the level of automation, enabling fastest time to prototype and its congruency with emulation. Designs that run in emulation through the Cadence multi-fabric compiler will come up in weeks or even days when mapped to the Protium S1 platform. No RTL changes are required, partitioning and memory compilation is fully automated, and FPGA place and route is fully integrated. Bottom line, the manual efforts that traditional prototyping requires are fully automated and users typically see designs running at 3MHz to 5MHz out of the box, around five times faster than in emulation. But that's not where performance stops. The platform is scalable from 3MHz to 100MHz, from fully automatic to fully manual. It's automated front-end allows further optimization of clocking, critical paths and memory ports. Together with higher effort place and

route, this easily extends the speed range to 10MHz. To extend performance even further, the Protium S1 hardware can replace home-grown FPGA systems easily and, with manual optimization, allows them to get to tens of MHz performance. For software development, the platform offers very specific, unique capabilities not found in any other FPGA-based prototyping systems. Memory upload and download together with capabilities to freely start and stop the design, and force and release signals, enable advanced software debug. It also has a transaction interface that allows host-based software to be connected. It is scalable from single-board designs to two, four and eight FPGA configurations scaling from 25 million gates to 200 million gates per chassis. It features multi-user support and scales to multiple chassis to 600MG and more.

To cover a wide range of design sizes and accommodate different interface requirements, the Protium S1 platform comes in multiple hardware configurations, ranging from two FPGAs per system to eight FPGAs per system. All FPGAs are Xilinx Ultrascale XCVU440 devices providing, depending on the design to be prototyped, up to 25 million ASIC gates capacity and up to 88Mb of embedded memory per FPGA. The boards are mounted in a custom chassis with power supply, cooling and all necessary interfaces and cabling included. In addition to being equipped with a variety of on-board interfaces, the platform is complemented by a comprehensive portfolio of daughtercards and is also fully compatible with Cadence's family of SpeedBridge adapters, enabling a smooth transition from an emulation environment to a FPGA-based prototyping environment. The resulting FPGA-based prototyping portfolio is shown in figure 2. As part of the Cadence Verification Suite, FPGA-based prototyping just became easier with the Protium S1 platform. ■

Product News

Hall-Stand 4891

SEGGER: secure data exchange for embedded IoT devices using Dropbox

SEGGER announces a Dropbox client as part of their Internet of Things offerings. IoT devices have different communication needs. The new Dropbox client is ideal for projects which need to manage various amounts of data. This is accomplished by loading or storing files, even large data files, in a secure and reliable way using the public Dropbox API. Typical use cases include; firmware updates, log files, and for that matter, any information shared between different IoT devices and/or a central server.

News ID 4891

Hall-Stand 4875

Artila: FreeRTOS programmable remote I/O module

Artila Electronics release RIO-2010PG, the new FreeRTOS programmable remote I/O module. RIO-2010PG is powered by a 32-bit NXP LPC1768 100MHz ARM Cortex M3 processor which is equipped with 64KB SRAM 512KB Flash and FreeRTOS operating system. The industrial I/O of RIO-2010PG features one 10/100 MHz Ethernet port, one full modem RS-232, one isolated RS-485, eight channels of relay, 16 photo-isolated digital inputs and one 1-wire interface for temperature or humidity sensors.

News ID 4875

Hall-Stand 4932

ANSYS: simulation platform transforms product development and sparks innovation

Engineers, from the novice to the highly experienced, can now use engineering simulation software across the entire product lifecycle with the newly released ANSYS 18. This feature-rich release expands the boundaries of simulation upfront in the development process to digital exploration as well as downstream with digital twins, expanding simulation to the operations and maintenance of products.

News ID 4932

Hall-Stand 4889

Atlantik Elektronik: ConnectCore for i.MX6UL development and rapid prototyping

The ConnectCore for i.MX6UL SBC provides a universal and powerful platform for your smart connected devices in the form of a compact and flexible development kit for rapid prototyping. The ConnectCore for i.MX6UL Development Kit delivers a complete off-the-shelf Single Board Computer (SBC) built on the ConnectCore 6UL System-on-Module. The new Digi ConnectCore for i.MX6UL SBC is now available at Atlantik Elektronik.

News ID 4889

Hall-Stand 4946

PLS: UDE 4.8 simplifies trace analysis and evaluation of runtime behavior

PLS Programmierbare Logik & Systeme is presenting version 4.8 of its Universal Debug Engine for the first time at embedded world 2017. The UDE 4.8 features a large number of completely new and improved functions for analysis of very large amounts of trace data and for evaluating runtime behavior of real-time operating systems. The enhanced trace analysis of the UDE 4.8 now allows developers to search even faster through this very large amount of trace data. The 'Find all' function of UDE 4.8 searches not only for single events such as, for example, function entries or accesses to specific memory locations, but also for entire sequences of events in one single search run through the complete dataset.

News ID 4946

Security has more dimensions than just crypto algorithms!

By Marco Blume, Wibu-Systems

This article highlights CodeMeter Embedded 2.0 which comes packed with new features and abilities and can now access key storage with multiple processes. This might not sound like a great deal, and it has long been standard on desktop PCs, but it has to date not been possible on embedded systems.

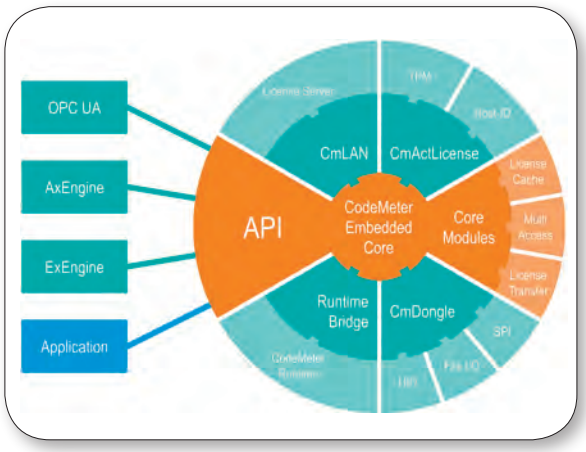

Figure 1. CmEmbedded 2.0 includes now License Core which allows the integration of additional software components.

■ CodeMeter is far more than dongles for protecting software – CodeMeter is a complete solution, an entire universe of components from development tools to the integration of license creation into sales processes. Crypto libraries are just one part of the edifice. Many providers have come up with AES encrypted software code and integrity protection mechanisms based on elliptical curves (ECC). They are on the right track, and their technology is absolutely state of the art. CodeMeter is also based on it, but it does not stop there: the best encryption algorithms are immediately useless whenever the keys are not stored securely as well. And even if the keys are kept secure in a Trusted Platform Module (TPM), the problem is only half-solved. How do they get there? Can they still be updated after 10 years of service in the field or when the system is offline? Is the system still as secure as it should be? How can developers and clients keep track of their keys on all their devices? All these vital questions can be solved, as we will prove.

On top of protecting the invaluable know-how of developers, businesses have another important goal: to earn money with their devices or software. This monetization essentially means enforcing some form of licensing system. To do so, the software needs to be protected from illicit use, closing the circle and taking us back to encryption. Beyond preventing or allowing a software to be run (i.e. decrypting it), licenses can have many other features. They might limit the duration or number of times a software can be used. They can contain protected data, like operating parameters. Additional licenses can be used to unlock add-on features. Of course, licenses can be tied to a specific system to stop illicit copying. Everything works together – CodeMeter users get know-how protection and the opportunity to introduce versatile licensing concepts that might open up complete new business opportunities.

Scalability to serve all systems

CodeMeter scales from servers and workstations, via the entire universe of different embedded systems, down to the tiniest microcontrollers. A fast-growing type of computer is entering our lives, not the old box on people desks, but tiny computers built – embedded – into other devices or machines. CodeMeter Embedded (CmE) is the CodeMeter version for embedded systems like industrial controllers or mobile devices, i.e. for anything that is neither a traditional PC or laptop or a server in a rack. This means that CmE serves the market that uses 98% of all computer processors produced today. At the lower end of the scale, where FPGAs and microcontrollers start, the baton is passed on to CodeMeter µEmbedded. It is designed with even smaller systems in mind, while staying fully compatible with the licenses and infrastructure of the greater CodeMeter universe.

CmE has been around for several years. As the needs and expectations of users have increased and the systems have become more powerful in terms of computational power and storage space, the feature set has also evolved. It runs on all major platforms like x86, ARM, and PPC and virtually on any operating system. The automatic software protection features of AxProtector are available for Linux, Android, and Windows. Integration with VxWorks does even further, and the features can be accessed directly with the tools of Wind River and used in a secure boot process in cooperation with an UEFI. For other systems, CmE is available in source code, allowing clients to use it in special real-time operating systems or bare-metal implementations (directly on the hardware without any operating systems).

Secure Key Storage

Put simply, CmE is the agent working between protected applications and key storage. It behaves not unlike a driver:it gives the protected program an API for executing cryptographic operations, updating licenses, or storing data securely. On the key storage side, it supports different types of storage and handles them transparently for the API – again

Figure 2. CME 2.0 can be used on a variety of hardware components.

reminding us of how drivers work. The crypto API is only part of the solution. The keys need to be securely stored as well to achieve true security. In the CodeMeter universe, this can be done on a piece of hardware or in a file. Hardware in this case means a secure element – a specially customized chip that stores all the secret information well away from prying eyes. Cryptographic operations are executed on the chip itself, and the keys never need to leave their secure home. Only the results of the operation are returned to the outside world. This is the strongest form of security, and it is used in a similar form in cell phone SIMs or TPM chips. The hardware secure element can come in many shapes and sizes, from USB dongles or microSD cards to ASICs that communicate with the system via an SPI interface. If hardware is not an option, the software solution, a CmActLicense, provides the same functions as a dongle. The only difference is that the cryptographic operations will not happen in a dedicated piece of hardware, but on the protected system itself.

The CmActLicense file remains protected and is tied to several unique traits and markers of the host system to prevent it from being copied. Another secure option for storing keys is to rely on networked CodeMeter servers that can supply several devices with the licenses they need, without each device having an individual license installed. The necessary communication is again protected with changing transport keys to prevent wiretapping and manipulation. For the CodeMeter API, all of these variants are fully transparent and can be mixed and matched as needed, since all key storage options are designed to be compatible with each other.

License Lifecycle Management

The second challenge remains: how to get the keys and licenses into their container? This is again a more complex operation than a simple key transfer – the licenses need to be sold, created, and managed correctly. This is made easy with a simple tool called CodeMeter License Central. Like the other solutions, it has been developed to scale perfectly from handling single user licenses on a few individual computers, with licenses created manually, to small-scale online shop operations and even large corporate ERP systems, with dependable high-availability hosting in any given data center. In CodeMeter License Central, individual licenses are created for each target device. The update files that contain these licenses can be transferred by virtually any medium, as they are encrypted and tamper-proof. They can only be decrypted in the specific target container, be it a dongle or a CmActLicense. The same mechanism can be used to license additional add-on functions after the original sale or to renew licenses limited to a specific duration or volume. In the same manner, trial licenses can be upgraded to full licenses.

Automatic Encryption Processes

AES encryption is known to be secure. If anything is cracked, the problem usually lies in a poor implementation, not in the AES algorithm itself. Wibu-Systems want its customers to concentrate on what they do best. That is why the CodeMeter Protection Suite gives them a powerful tool to take over the task of encrypting and protecting applications or libraries. Developers can protect the software they have completed without knowing anything about encryption and directly activate the licensing system. CodeMeter Protection Suite does the heavy lifting for them: encrypting the software, integrating the decryption tools, and adding the meta-data – a great tool and the product of over a quarter of a century of experience in software protection.

CodeMeter Embedded 2.0 vs CmE 1.0

Why choose CodeMeter Embedded 2.0? Because it comes packed with new features and abilities. CodeMeter Embedded can now access key storage with multiple processes. This might not sound like a great deal, and it has long been standard on desktop PCs, but it has to date not been possible on embedded systems. The code was meant to be as lightweight as possible, it should not have any services running, and CmE was to be a directly integrated part of the encrypted program. The advent of more powerful embedded systems, however, brings new capabilities, which means that several protected programs and processes could share access to a single CmContainer. Changes to the inner workings have made the licensing system more flexible and effective. We are calling it License Core. The new functions allow the integration of additional software components. Based on License Core, an OPC UA stack can now be used with CodeMeter as key storage. Updating CodeMeter licenses is now also possible via the OPC UA protocol.

By relying more on secure elements in chip format that are built into new hardware already in the design stage, the SPI interface is integrated to handle communication. This avoids the detour via the USB stack, saves energy, and accesses the chip directly. CmE 2.0 is putting down the groundwork for the license transfer feature introduced in the desktop version in 2016. The new functions will be supported in embedded systems as well as new functions and features are being rolled out. CmE will never be a run-of-the-mill product – many of its features are designed to be modular. The mission has been and still is to produce the most compact software possible.

That is why we first ask new clients about their target system and use case before we produce the right package for them. This means that CmE can scale to match the client's needs. It is never a monolithic block of deadweight on the system resources. Wibu-Systems today offers a licensing and protection solution for almost any device that contains a processor. We are supporting developers with the implementation and sales professionals with the management of their licenses. ■

Analyzing real-time operating systems using trace techniques

By Jens Braunes, PLS

Code and data trace have absolutely no influence on the real-time behaviour of a system. This advantage can be used extensively when analyzing the runtime behaviour of real-time operating systems.

■ From experience, system observation is among the most critical aspects of testing and analyzing embedded software, because many factors may adversely affect runtime behaviour. One of them arises already during the build process: do we compile the application for debugging and test purposes or for production enabling additional optimizations for code size or speed? Profiling, which is frequently applied in order to optimize performance, also has a considerable influence. Quite often, the application needs to be instrumented for measuring the runtimes of functions and tasks. The added test code will have a small, but measurable influence on the application runtime behaviour. Even the memory layout may change due to the additional code, so that the influence cannot be neglected anymore. Among other variants of high-level system observation, monitoring is frequently used especially in the context of real-time operating systems. For this purpose, the monitor checks for RTOS state modifications that are typically mapped to specific memory locations. Triggered by interrupts which occur on task switches (among others), the monitor records the event for later analysis.

Those engineers who want to go without any instrumentation of the program code and without any monitors for analyzing the runtime behaviour of a real-time operating system (RTOS), should select a microcontroller providing appropriate trace hardware including a high-bandwidth trace interface. MCU families offering suitable features are now available from almost all semiconductor vendors. Today, tracing capabilities are often a show-stopper for the selection of a platform. Tracing can be used to observe modifications of system states without influencing the system real-time behaviour. Depending on the controller manufacturer and the implemented trace architecture, there will be several trace modes which might be useful to analyze the runtime behaviour of an operating system. However, before going deeper into that, it is necessary to point out what information is required actually in the specific case.

In RTOS-controlled applications, the runtime of each specific task plays a significant role for evaluating execution performance. For instance, the task execution time directly indicates whether the system is running at a reasonable workload or in an overload situation. As another important factor, it is interesting to know how often tasks are interrupted by other tasks and how long tasks can operate without being interrupted. Due to the overhead incurred by task switches, this often bears significant potential for optimization. The same applies for the interrupt load which gives an indication of how often and how long task execution is halted by interrupts.

Obviously, code tracing is the method of choice for obtaining the required runtime information for all necessary measurements as simply and quickly as possible. Basically, tracing consists of recording the addresses of executed branches (more specifically of the branch targets) as well as any deviations from the regular, sequential execution flow as a result of calls, returns and interrupts. For task analysis, code tracing can be used to detect task switches that manifest themselves by multiple function calls and returns (including the scheduler), thus causing disruptions of the sequential execution flow.

However, the code trace is much too large for this purpose because trace data are also recorded for the execution of in-function code, including if-then-else constructs or loops. This increases the size of the trace and represents an unnecessary load for the on-chip or external trace memory. Infineon uses an interesting approach to avoid this dilemma. The so-called Compact Function Trace (CFT) of this company records only function calls and returns. The program flow within functions is not captured by the trace at all. Of course, that reduces dramatically the amount of trace data and saves trace memory. Data trace is another option complementing code tracing. Typically, RTOSs have their own management structures

ORTI - C:\OSEK Awareness\orti_tc27x\orti_tc27x.ort

Name	Value
TriCore_27x	
vs_SMP_NUMCPU	0x3
RUNNINGTASK	
[0]	/M_Core1
[1]	/M_Core1
[2]	/M_Core1
RUNNIN	
SERVICI	
RUNNINGTASKPRIORITY	
CURRENTAPPMODE	OSDEFAULTAPPMODE
TASK_PWM_Core1	
TASK_1MS_Core1	
PRIORITY	"osTcbActualPrio[1]" not a valid expr
vs_HomePriority	11
STATE	READY
STACK	osTaskStackCore1_1
vs_Schedule	FULL-Preempt
vs_TaskType	EXTENDED
vs_WaitMask	0x0
vs_EventFlag	0x0
CURRENTACTIVATIONS	0x1
vs_MaxActivations	1
CONTEXT	Context_TASK_1MS_Core1
TASK_IoHwAb_Core1	
TASK_Get_Check_Tempt_Core0	
TASK_BSW_MainFunction_Core0	
PRIORITY	"osTcbActualPrio[4]" not a valid expr
vs_HomePriority	2
STATE	READY
STACK	osTaskStackCore0_1
vs_Schedule	FULL-Preempt
vs_TaskType	EXTENDED
vs_WaitMask	0x0

Refresh Values F11
Refresh ORTI File
✓ Trace RUNNINGTASK
Apply Trace Configuration

Figure 1. Processed ORTI file

located in MCU memory reflecting the current system state and the task scheduling. This includes active tasks and information about the interrupt processing. Changes of the system state, including task switches, thus always trigger a write access to these management structures. Some trace implementations allow data trace that can be used to record these write accesses. That enables a precise tracking and analysis of state changes of the operating system. However, data tracing is quite expensive in terms of memory space. Without additional measures, all memory accesses of the application would be recorded. Similar to code tracing, this is undesirable here as well. Appropriate filters must therefore be applied to ensure that recording is exclusively restricted to relevant accesses targeting the management structures. Regardless of which trace mode is used for obtaining runtime information, unique, uniform timestamps are required here. Otherwise, it would be virtually impossible to get precise measurements and useable results.

How an analysis using trace data will work in the field can be demonstrated with an example of an OSEK-compliant operating system. OSEK is the abbreviation for "Offene Systeme und deren Schnittstellen für die Elektronik in Kraftfahrzeugen" (open systems and their interfaces for automotive electronics). It designates a standard for implementing real-time operating systems particularly in the automotive environment. The well-known Autosar approach represents an advanced development and reuse of the OSEK specification. The OSEK also defines an interface format called ORTI (OSEK Runtime Interface) for the communication of the operating system with analysis tools and debuggers. In a text file (the so-called ORTI file) this format describes all relevant internal operating system data, facilitating their use by tools and their visualization for users.

The debugger - the Universal Debug Engine (UDE) from PLS in this case - extracts the data structure from the ORTI file containing for example the current task (figure 1). This information is used for configuring the trace filter, ensuring that only data trace is recorded

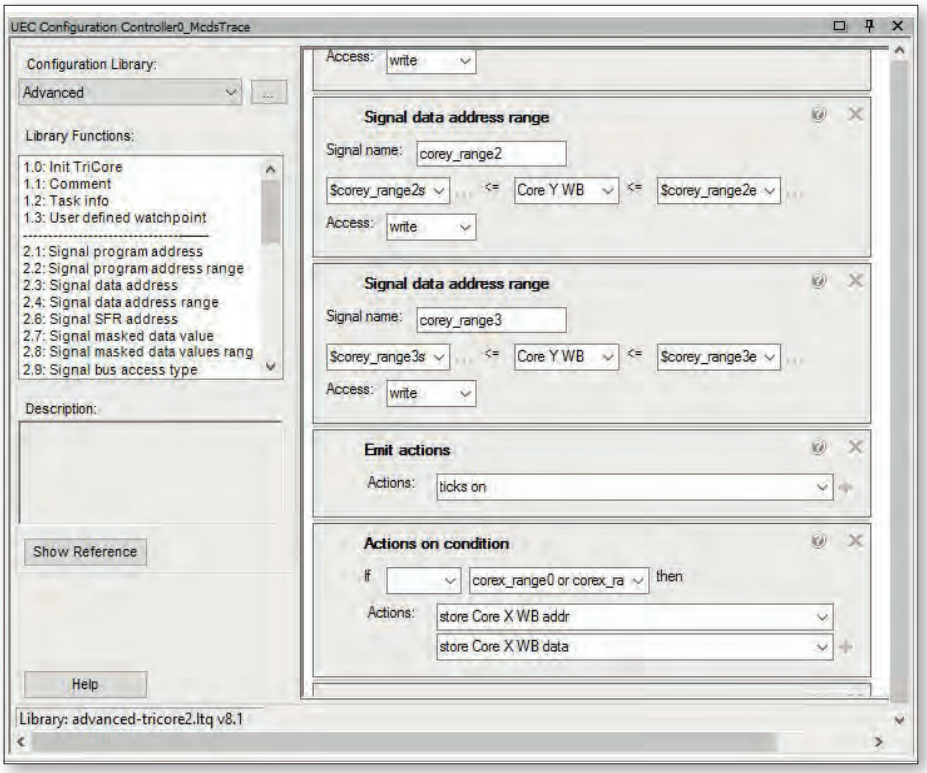

Figure 2. Configuration for a data trace based on ORTI data

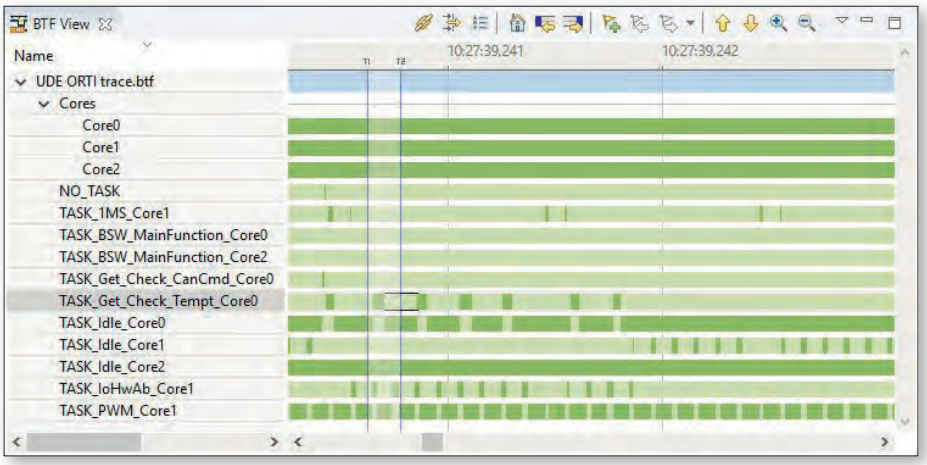

Figure 3. Visualization of operating-system tasks in Eclipse Trace Compass

actually relevant for task analyses (figure 2). In the next step, the trace is preprocessed to allow a proper visualization of the sequences of executed operating-system tasks including their precise timing. Specialized tools like Eclipse Trace Compass (www.tracecompass.org) can be used afterwards for sophisticated visualizations and additional analyses (figure 3).

For this purpose, the UDE provides an export function based on the BTF format (Best Trace Format) defined by Timing Architects, Inc. BTF, which was specifically developed as an exchange format for event traces, is used in many simulation, profiling and trace analysis tools. As demonstrated by this example, a combined tool solution consisting of a debugger and hardware tracing is now a genuine alternative to collecting runtime information using instrumentation or special monitors. Even though hardware trace will often operate too finely-grained - it will typically record the entire control flow including in-function instructions and will thus be quite memory-hungry - developments like CFT or the use of ORTI for data tracing can quickly compensate for this disadvantage.

All in all, tracing provides a method for collecting runtime information in a fully reactionless and highly precise manner. And not least, it can also be used for examinations at the task and operating-system level. ■

Product News

Hall-Stand 4929

Keysight to address IoT, digital, RF test challenges at embedded world 2017

Keysight Technologies will exhibit a number of new products at embedded world 2017. Keysight's technical experts and application engineers will be available at the event to demonstrate the company's latest design, verification and test solutions. The solutions are focused on the general electronics, automotive, communications, education and energy industries and were designed to help embedded engineers solve today's toughest test challenges. This includes IoT projects, digital and RF communication, and (low) power requirements. Demonstrations at the booth will include the following comprehensive set of hardware and software-based solutions:

Entire range of low to high-end oscilloscopes – showing the latest solutions for debugging low- and high-speed serial busses and the testing of communication signals. This includes USB-C, NFC and wireless charging application support.

RF and wireless communication solutions – including the EMI/EMC aspects of the IEEE 802.11 protocol (WiFi), Bluetooth and ZigBee channels. Simulating and testing these wireless channels, which are used in a wide range of embedded applications, have become a key, and essential step in the embedded design process.

New low power and ultra-low power measurement solutions – which are relevant to embedded designers working on IoT, semiconductor, automotive and wireless applications. This includes the CX3300A, industry's first 100pA dynamic current analyzer with 14/16-bit resolution and up to 200 MHz bandwidth and up to 1 GSa/s sampling rate.

New hardware and software tools for signal integrity analysis and verification – including Keysight's new simulation tools SIPro and PIPro, allowing designers to investigate signal integrity and power integrity at the design stage. This also includes a new test solution for evaluating thermal issues on the board as well as measuring TDR. New IoT measurement tools – including a teaching kit with courseware to assist universities in building competencies to match industry needs in product design, validation and testing, and manufacturing.

General purpose instruments – offering a comprehensive portfolio of bench and PXI instruments for precise and reliable measurement needs.

News ID 4929

Product News

Hall-Stand 4900

Wibu-Systems: security extension of CodeMeter integrates with OPC UA standard

At Embedded World 2017, Wibu-Systems will prove how CodeMeter's protection, licensing, and security technology is serving the embedded system community with the most advanced, modular, and neatly designed architecture available today. CodeMeter provides cybersecurity capabilities to intelligent device manufacturers who want to safeguard firmware upgrades and updates, to software publishers that need to protect their intellectual property from counterfeiting and reverse engineering, and to users who want to be sure the software they run is genuine and has not been tampered with. CodeMeter also includes the entitlement aspect, from the creation of licenses, to their online or offline delivery, transfer, or remote management. In doing so, Wibu-Systems supports the migration to a dynamic industrial economy where new license-based business models create additional revenue and responsive pricing models for an unlimited access to the market for all vendors and users. Wibu-Systems will show the first samples of CmStick/CM, the ultra-compact, simple, chic and robust USB stick with full metal housing, high performance flash mass memory, USB 3.1 interface, and a System in Package module for the USB connector to withstand intense vibrations, high humidity, and temperature fluctuations. In order to meet the most stringent durability and reliability standards for data integrity, this model, just like all other variants in the CodeMeter stick family, will embed the state of the art in its bill of material: the avant-garde SLM97 security controller from Infineon and Hyperstone's U9 flash controller and their patented hymap firmware. It will also come with CodeMeter API for the read and write configuration of a secure area in the flash memory.

News ID 4900

Hall-Stand 4935

Logic Technology: Reliance Edge 2.0 improves performance and flexibility

With Reliance Edge 2.0, Datalight delivers improved performance for flash-based devices with the introduction of discards in the commercial version, and gives more flexibility with the addition of support for Linux as a host environment. Reliance Edge is available under a GPLv2 license or can be licensed commercially. For the open source version, support for Linux as a host environment has been added, making it much easier to use Reliance Edge with projects being developed using Linux. The bug fixes in this release correct directory and mount operation issues.

News ID 4935

Hall-Stand 4917

Cadence: fast path to PCB power integrity signoff

Cadence Design Systems announced availability of the Sigrity 2017 technology portfolio, which introduces several key features specifically designed to speed up PCB power and signal integrity signoff. Among the features included in the newest version of the Sigrity portfolio are the Allegro PowerTree topology viewer and editor, which enable designers to quickly assess power delivery decisions early in the design cycle. The latest release of Sigrity also includes a PCI Express 4.0 compliance kit for checking signal integrity compliance with the latest PCIe specification when it is certified later this year.

News ID 4917

What you say is what you get: an Ada story

This article is contributed by AdaCore

Over the past 30+ years, Ada technology has matured into a unique toolset allowing programmers to achieve software reliability at a very affordable cost. It's available for small microcontrollers such as ARM Cortex-M, large x86 64 hosted systems, and many things in between. Time to give it a try?

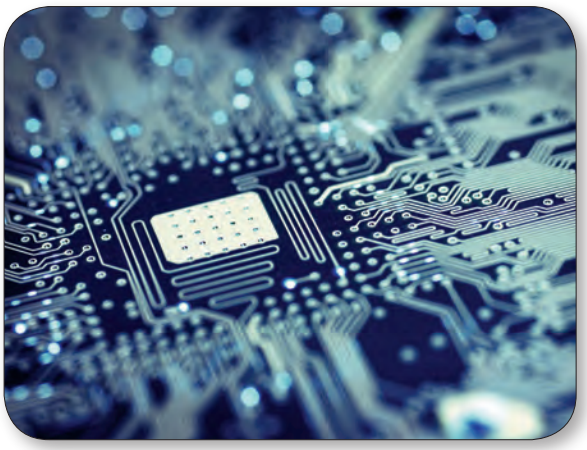

■ At the risk of opening with platitudes, software development has become one of the most central engineering disciplines. Our lives are literally governed by software in quantities that defy imagination. As a matter of fact, our lives actually rely on this software. Drones, cars, medical devices, the list of things whose software can make the difference between life and death is growing by the day. How to avoid making this a list of threats? Truth is, software engineering is not about writing code. It's about specifying intent, implementing what's intended, and verifying that what's implemented corresponds to what's intended. And to do that, all environments and all languages are not equal.

Ada is a language that was first defined in the 80s, around the same time as C++. Unlike C++ though, which was designed as an evolution of C with extra expressivity capabilities, Ada was aimed at supporting the needs of high-criticality embedded software development. This led to a series of unique technical choices setting the language on a very different path. As a result, Ada tends to force programmers to be much more explicit when they write code, as the compiler will prefer to avoid compiling rather than guessing. Conversely, as opposed to other languages, Ada allows programmers to formally express intent that would be otherwise buried into comments.

This has two interesting consequences. First, many programming errors that would be difficult to avoid can be mitigated or eliminated. Second, it's actually possible to go much further in making explicit expectations and intent at the software levels and either automatically enforce it or verify correct implementation. Here's an example of what that may look like at low level:

```
type Percent is delta 0.1 digits 4 range 0.0 .. 100.0;
type Status is (On, Off, Error);

type Sensor is record
   Power   : Percent;
   Enabled : Status;
end record
with Bit_Order                    => High_Order_First,
     Scalar_Storage_Order         => High_Order_First,
     Size                         => 32
     Dynamic_Predicate            =>
     (if Sensor.Enabled = Off then Sensor.Power = 0.0);

for Sensor use record
   Power at 0 range 0 .. 15;
   Enabled at 0 range 30 .. 31;
end record;
```

This says quite a lot about a low-level data structure which is possibly used to map onto a device, or a data connection. It's a sensor structure composed of Percent, which is a fixed-point value with a precision of 0.1, from 0 to 100 of 4 precision digits, stored in the first half-word of a 32 bits big-endian structure, and a flag that can only take 3 values, stored in the last 2 bits of that same structure. From this description, the compiler will actually verify that this representation can be implemented and if it does, guarantee that this specification will indeed be respected. Developers familiar with development of low-level layers can already see macro and bitwise operations headaches vanishing thanks to these features. In addition to this structural information,

the type is associated with a constraint (or dynamic predicate) that states that when the flag Enabled is Off, Power has to be equal to 0. With assertion checking enabled, the compiler will generate consistency checks at appropriate places in the code, allowing the programmer to identify inconsistencies early.

On the opposite side of the spectrum, specification can be used to express functional constraints. The most common one is probably the precondition. The idea is that an assertion can be defined to be true when calling a function (or subprogram). A postcondition will specify the output of a function - its behaviour. Take for example:

```
procedure Swap (A : in out An_
Array; I1 : in Integer; I2 :
in Integer)
   with Pre => I1 in A'Range
and I2 in A'Range,
        Post => A(I1)'Old =
A(I2) and A(I2)'Old = A(I1);
```

The above defines a simple swap procedure on two elements on the array. One may expect that the indexes are within the range of the array. To avoid out of range array access, it's tempting to write some defensive code within swap to detect the errant case and take appropriate measures - raise an exception or cancel the operation. But it should really be the responsibility of the caller to correctly call swap with two valid indexes. The precondition of this code allows the programmer to specify exactly that, thus removing the need for defensive code. It literally says I1 has to be within the range of A, and so does I2. On top of this, we also specify the outcome of the subprogram, that is the old value of A(I1) now equals to A(I2) and vice versa. So not only is there verification on the way in, to check that the code is properly called, but there's also verification at the point of return, to check that the operation is doing what is expected. Also, note the parameter mode in out on the array: that specifies that Swap will modify the array, as opposed to I1 and I2 which are in (i.e., only input).

Preconditions, postconditions and other kinds of functional contracts (such as dynamic predicates as shown earlier) can be verified in many different ways. The obvious one is to ask the compiler to generate actual checks. This can be done selectively. All assertions could be kept during testing for example, and only a minimal amount kept after deployment. Even when inactive, these contracts can have merit just from the fact that they're written and accepted by the compiler – i.e. they refer to actual entities and consistent operations. And they also serve as useful information to the human reader.

There's however another way to do this verification that doesn't even involve executing the code: static analysis. Better yet, formal proof. The SPARK language is a subset of the Ada language that is fit for static verification. Within the subset, any potential error will be analyzed and those that may happen will be reported. Every contract will be looked at taking into account all possible values, and those that might not be valid will be reported. In particular, SPARK can guarantee that every single call always respects the preconditions, and that every single subprogram always respects its postcondition.

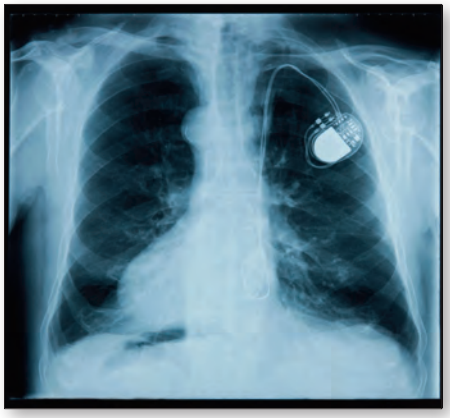

And there's more with Ada. Class hierarchies with verification that the behaviour defined at the root level is consistent with the behaviour of subclasses. Concurrency models with guarantees of absence of deadlocks. Specification of global variables - and more generally global states. Physical dimension consistency checking. Integer overflow elimination. The list goes on.

So who is using Ada today? While the language has a long-standing history in aerospace and defense, it's now frequently being adopted outside its original area by a growing number of companies in other embedded domains. Some are in the very enviable situation of having projects to start from scratch. Adopting Ada or SPARK is very easy then, it's merely a question of training and there's nothing fundamentally difficult for a typical embedded developer. Most of those companies, however, have a significant investment in C or C++ code bases - or rely on components developed in those languages. So how do they migrate to Ada? Surely, bearing the costs of rewriting the whole code base offsets any benefits they may get.

Luckily, this is never a requirement. Ada is extremely amenable to language interfacing. As a matter of fact, often C and Ada compilers share most of their code (at least the part responsible for optimizing and generating assembly) making mixing those two an easy task. As an example, the swap function described already could actually come from a C library:

```
void swap (int * a, int i1,
int i2) {
   int tmp = a[i1];
   a[i1] = a[i2];
   a[i2] = tmp;
}
```

All we need to do is to tell Ada that this implementation is actually in C, specifying that it's imported, of a C convention, and linked to a symbol swap:

```
procedure Swap (A : in out An_
Array; I1 : in Integer; I2 :
in Integer)
   with Pre => I1 in A'Range
and I2 in A'Range,
        Post => A'Old (I1) =
A(I2) and A'Old (I2) = A(I1),
        Import,
        Convention => C,
        External_Name =>
"swap";
```

Hey, isn't that adding contracts to a C function? Note that although this interfacing code can be manually written, there are also binding generators available that would automatically take care of most of the burden (with the exception of addition of contracts). At the end of the day, only specific software components will be re-written in Ada, most legacy code can be kept. New modules developed in Ada can then be integrated with this legacy code. What you say is what you get. Software engineering is not about writing code, it's about specifying intent, implementing it and ensuring that what has been implemented corresponds to what has been specified, with as little effort as possible. ■

From CE to IoT Core: Microsoft and the rise of real-time applications

By Daron Underwood, IntervalZero

This article shows how a combination of low-powered, small form factor Intel Atom processors, Microsoft's Windows 10 OS for IoT devices, and a real-time platform like IntervalZero's RTX6 can provide a solution for the legacy of Windows Compact (CE).

Windows 10 IoT is suited for various applications.

As the Internet of Things (IoT) continues to grow in popularity and need, the systems that enable it must evolve as well. Microsoft's transition to a one-Windows world certainly benefits IoT systems, but is that enough? To meet the real-world requirements of this new ecosystem, organizations that build IoT apps need a real-time operating system (RTOS). A real-time deterministic solution ensures that your application's performance is not only maintained but enhanced, without losing memory or the capability to run other applications without slowing down your system.

Although Microsoft has made no direct comment regarding the future of Windows Compact (CE), it is clear that CE is giving way to Windows 10 IoT Core. The interesting thing with CE is that it had two distinct uses: products that needed real-time processing and those that did not, or at least did not need tightly bounded determinism. In fact, in the beginning, CE was a simple general-purpose operating system that was different than Windows, but provided a similar programming interface (Win32 subset) as well as a similar overall user look and feel. The original target for CE was small single-purpose devices.

However, once CE went into use, users quickly recognized that in many cases, the devices that would benefit from using CE also had a real-time, deterministic requirement; for example, an environmental control system for a building or a controller for brewing systems. Microsoft recognized this need for real-time processing in the OS and decided to work with industry experts, including VenturCom (now IntervalZero) to provide it as an inherent part of the product. So, in version 2.12, Microsoft introduced a real-time scheduler into the Windows CE kernel. Windows 10 IoT Core must now address the same two use-cases as its predecessor. It can certainly meet the non-real-time needs with flying colours, but in spite of the real-time priority class, the other is a real problem for Microsoft. How can they keep real-time users in the boat and attract new customers?

Windows CE ultimately failed to capture a huge market share because of its disjointed relationship with standard Windows. While there was a fairly good amount of overlap at the top of CE and bottom of WES, there were also some pretty big gaps. For example, you may have needed some amount of determinism but really needed the full power and ease of use of standard Windows. Additionally, while Windows CE was similar, it was different enough in user interaction and development tools to be distinctly not the same as standard Windows. Much of this issue stemmed from the fact that CE is a distinctly different kernel and subset of the programming interfaces. This was a big turn-off to many companies. Although developers could use a good amount of code for each system, there still had to be code sections that were specific for CE platforms and standard Windows platforms, including WES systems.

In contrast, Microsoft has refactored Windows 10 in such a way that it can build the various OS products, including IoT versions, from a core set of functionality which scales up to a given device/platform. This was a vital effort in providing a truly scalable OS platform. The result is known as OneCore. OneCore exists at the bottom of all Windows 10 platforms, including IoT SKUs. Out of it comes IoT Core, which can be built to run on x86, x64 and ARM devices. IoT Enterprise, like WES of the past, is an x86/x64 only OS. In fact, IoT Enterprise is exactly the same product as the Windows 10 Enterprise LTSB (long term service branch) with a different licence model than the IT Enterprise SKU. This is actually great news for IoT/embedded developers, as it ensures changes to the OS, including updates, only occur at relatively long-spaced intervals, unlike the frequently updated commercial version. UWP (Universal Windows Platform) and .NET Core take OneCore to the next level by providing developer services and application programming interfaces that work across

Windows 10 IoT is a OneCore OS, App and Device platform including runtimes and frameworks.

platforms. This now addresses one of the bigger flaws with CE. If you design to .NET Core, your application will run on the family of IoT devices from Core to Enterprise. This empowers scalability. At the system level, universal drivers like UWP and .NET Core are a set of programming interfaces that, if adhered to during driver development, can produce drivers from the same source code that can run on x86, x64, and ARM systems. This is huge for device vendors that provide hardware that can work on many different systems.

It may not be binary compatibility, but for device drivers, write-once-target-many is a powerful tool. Microsoft has addressed almost all of the previous major challenges to bridge the gaps between Windows CE and WES. The last and most important is the real-time processing component. Windows 10 IoT is great as-is for many products, such as kiosks, ATMs, vending machines, etc. But there is a huge number of devices and systems out there that require at least a minimum amount of hard real-time processing. Beyond those more purpose-specific devices, the ability to create large-scale deterministic processing systems is now being realized. To replace rigid hardware systems with the power of Intel architecture, a solution such as the combination of Microsoft Windows 10 IoT and a real-time platform like RTX64 is both needed and available today. These systems push the envelope in terms of multi-function processing, from taking advantage of the Windows ecosystem to meeting hard real-time requirements on the same system at the same time. With the combination of low-powered, small form factor Intel Atom processors, Microsoft's consolidation of the Windows 10 OS for IoT devices,

Windows 10 IoT is real-time enabled by RTX64.

and a real-time platform like IntervalZero's RTX64, there is a straightforward solution to the CE legacy. In fact, it has evolved tremendously. Imagine the ability to develop single real-time solutions that can work across a full range of small devices up to enterprise-level systems without any binary modifications. The ability to build these systems up from a common foundation, using higher and higher functional components is extremely powerful. Imagine a Windows 10 IoT platform that adds a real-time capability, topped by functional modules such as Programmable Logic Controllers (PLCs), vision processing engines, and motion control. All added as individual plug-n-play like modules that are easily integrated by the platform to create a fully functional, flexible and scalable system. Such a platform would be the ultimate goal for organizations leveraging IoT, as it would allow for a completely integrated solution wrapped into a UX-driven software PLC that contains all that an organization needs for extensibility and scalability for implementing IoT solutions now and into the future. ■

Addressing IoT impact on software engineering

By Bill Graham, GrammaTech

Manufacturers need to carefully evaluate the cyber threats and the level of exposure of IoT devices. New levels of software integrity can only be achieved if teams can eliminate both accidental coding errors and intentional design-in vulnerabilities, through efficient analysis techniques suitable for the typical highly complex applications of today.

Figure 1. CodeSonar has been proven to provide the deepest static analysis, finding more critical defects than any other static analysis tool on the market

■ Powered by the forces of the cloud, connected endpoints, wireless technologies, and big data, the Internet of Things (IoT) evolution is forming a perfect storm for software engineering teams. This single, transformative force is bigger than anything in the history of tech industry, fueling an unparalleled consumer-oriented features race, expected to advance at an incredible rate over the next decade.

And why not? Vendors are racing to claim a piece of the predicted 8.9 trillion dollar IoT market by 2020, made up of more than 50 billion IoT devices spanning nearly all markets – automotive, energy/utilities, home appliance, consumer electronics, medical, education, manufacturing, and more. Although exciting to consumers and businesses alike, this race for IoT dominance also brings a significant dark side as embedded applications open communication to send and receive data between networks and other applications that may not always be known or friendly.

Current manufactures are still developing products using old and entrenched supply chain, engineering, and QA processes that weren't designed for the complexities of highly-connected smart devices today. Engineering teams are utilizing a progressively diverse set of suppliers and relying on third-party software to save time while trying to satisfy the business and market thirst for IoT demands. Unfortunately, many software development teams treat security in much the same way they have in the past, running only basic checks, if any, during their QA cycle.

This confluence of drivers – the lack of a security-first engineering philosophy, the increased use of third party software, and the continually growing time-to-market pressures from business executives complacent about IoT security – will continue to put software engineers in an ever-increasing tough spot, ripe for cyber criminals and nation states looking to exploit these connected devices and networks. These software vulnerabilities have already put consumer safety and privacy at risk, increasing corporate liabilities, eroding trust, and in some cases, shutting down critical public and industry services.

The fact of the matter is that today smart devices are anything but smart. One recent study found that 70% of the top 10 IoT smart devices are vulnerable to exploitation. The daily onslaught of news reports regarding new devices, appliances, and systems that have been hacked, includes stories that are quite terrifying, such as hackers remotely taking control of an automobile through its wireless hot spot connection and successfully commanding brakes and other critical systems. And unlike typical server applications that are housed in secured facilities with restricted access and controlled network connections, IoT devices (and their applications) are easily compromised. Allowing unprecedented time and access to isolate, reverse-engineer and repeatedly stress test for weaknesses. The reality of this is that if software engineers are unwilling or unable to find the cracks in their applications, someone else will. Possibly with bad intent.

So how do we evolve manufacturing processes to better protect our next-generation IoT devices? First, it starts with a sound plan that includes next-generation software assurance and a security first methodology. Teams need to rethink how they deliver software quickly – with security, safety, and quality in mind from design to deployment. However, rethinking should not be restarting. For example, trying to train every developer in the tactics hackers might use to exploit their software is not only impractical, but very time consuming. To do this successfully, teams should leverage the best tools available that help them analyze the software they are developing looking for problems that IoT presents – including both in-house source and third-party binary code. As IoT applications become more feature-rich, with additional elements of internet-connectivity and device intelligence, the risks of

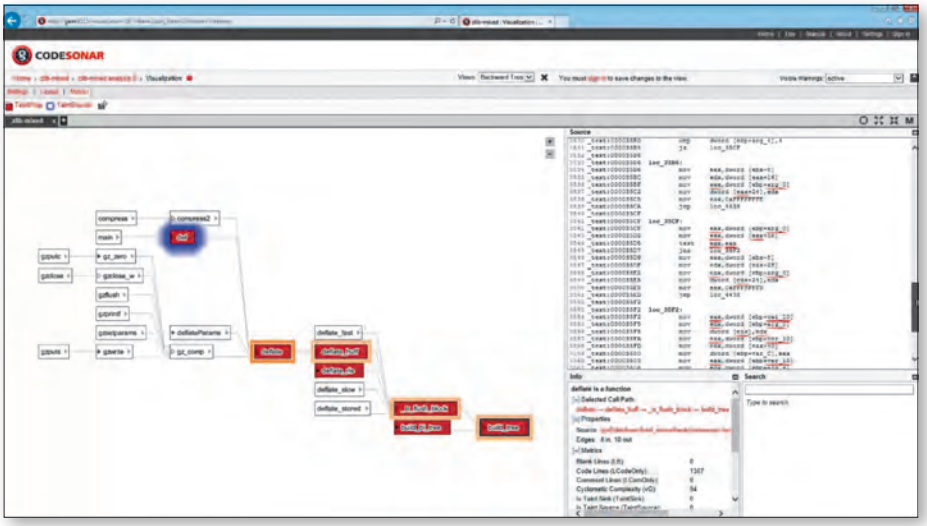

Figure 2. Taint Analysis allows to visualize notoriously hard-to-find tainted data pathways as an overlay on the code or superimposed on a graphical visualization of the program.

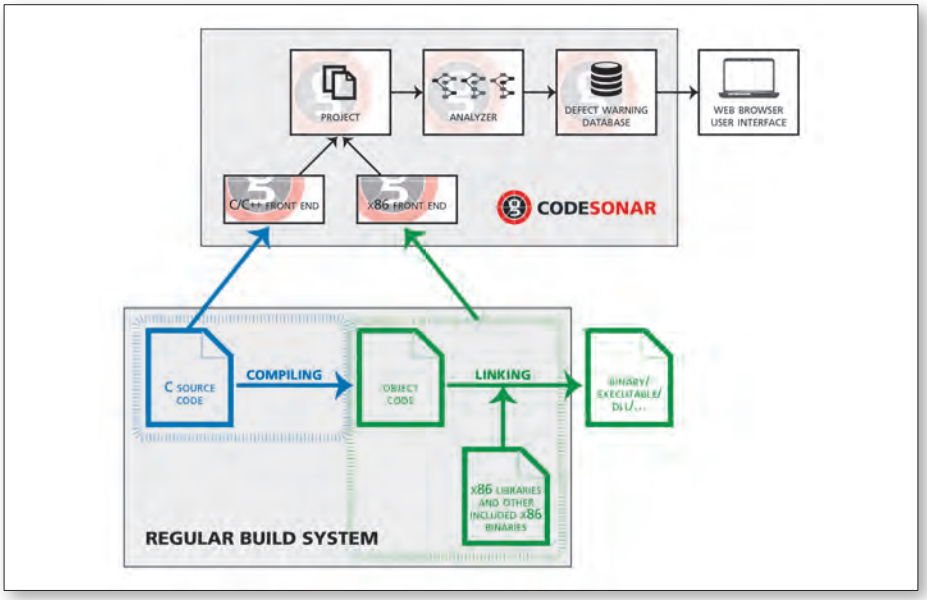

Figure 3. Due to the ever-rising amount of third-party code, strict check of third-party binaries is required. To realize this CodeSonar integrates binary code analysis.

built-in security vulnerabilities are increasing. Despite this trend, awareness of the risks associated with insecure code is still low among IoT developers and QA teams, and not a priority with most management teams.

Modern static-analysis tools are popular because they have proven to be effective, they are simple to introduce, and they can be used by development, QA, and security audit teams. Furthermore, in contrast to traditional dynamic testing, the code analyzed is never executed, so there is no additional test case development overhead and static analysis can be applied very early in the development process.

When programmers use static analysis as soon as code is written, bugs and security vulnerabilities can be found and eliminated even before the unit testing or integration testing phases begin. The earlier a defect is found, the cheaper it is to fix; this cost saving is a major advantage of automated static analysis.

Fortunately, static-analysis tools for source and binary have the ability to detect vulnerabilities before products are shipped, dramatically reducing security threats and corporate exposures that cost organizations several millions of dollars.

We've seen this numerous times in the recent news – with Toyota's unintended acceleration issue (estimated $3 billion in costs), in addition to the brand's first black eye; potential safety hazards with Jeep's recently-hacked Uconnect vulnerability, affecting over 470,000 vehicles; and the several SCADA systems recently hacked, most notably the Stuxnet exploitation, used to attack and destroy industrial equipment. It's simply unacceptable, and more importantly, easily avoidable, for development teams to not implement proper checks for security, reliability and hacking today. The added level of software assurance, which CodeSonar provides, can be easily deployed for the cost of a developer's morning coffee and a scone.

Over the last few years, third-party code has moved from a minor factor in software development to a dominant force in the industry. It is now used throughout software development in all applications, from highly sensitive government applications to security-intensive financial systems to safety-critical applications to consumer and mobile applications.

According to the latest report from VDC Research, the majority of software that runs on embedded devices is now developed by external sources, not in-house development teams. Some of this is open-source, but in embedded applications, nearly 30% of code is third-party commercial software – so the source is often unavailable. Such components include graphics toolkits, cryptography libraries, and communications middleware (network, USB, Bluetooth), which make up nearly 70% of the common embedded attack vectors.

GrammaTech, leveraging over 10 years of collaborative research, has developed a binary analysis capability to examine third-party code without requiring access to source code. This capability is fully integrated within our proven static analysis tool, CodeSonar, the first and only commercially-available binary analysis product.

CodeSonar's binary analysis technology provides developers with the ability to evaluate, check, and inspect third-party code, and provides businesses with more options within their supply chain, enabling them to utilize software from new, innovative companies that might not have an established reputation. When source code is available, you can use CodeSonar in mixed source/binary mode, analyzing your complete application.

The days of developing a standalone application are gone – the Internet of Everything has rapidly forced manufacturers to rethink how their products will support connected economy today, and changed the threat landscape forever. Nowadays reality is that there are educated attackers whose sole function is to break into IoT systems for many reasons, including fun, intellectual stimulation, profit, or worse, offensive attacks and terrorism. Today software development teams must adopt a robust secure design lifecycle, giving them the insights and capability to get it right first, to prevent these attackers from having a

chance at breaking in. A general rule of thumb for teams to follow involves an end-to-end threat assessment (from a third-party audit team), security optimized designs, and security-scanning tools, of source and binaries. The Internet of Things is a paradigm that is impacting our daily life – for good and bad – today. IoT software needs security by design; for this reason, it must be a business imperative. Manufacturers must carefully evaluate the cyber threats and the level of exposure of IoT devices, implementing all the necessary design checks and countermeasures to respond to an accelerating set of menaces. GrammaTech was founded 26 years ago, with a firmly-grounded purpose to help organizations develop tomorrow's software. Given the ever-increasing dependence of software in today's connected world, our experts are focusing on solving the most challenging software issues through a thorough portfolio of software and security assurance solutions. IoT is here, it is our responsibility to ensure our software is ready for it. ■

Product News

Hall-Stand 4914

Mouser and SparkFun bring open source hardware to global markets

Mouser Electronics announces a global distribution agreement with SparkFun Electronics as part of both companies' commitment to support the growth of the maker movement around the world. The agreement provides increased access to both Mouser's and SparkFun's extensive product ecosystems, enabling customers to easily shop for the right maker technology to support their projects. Mouser's catalog now incorporates over 500 SparkFun products, including SparkFun's signature Arduino Pro, RedBoard, and LilyPad tools, empowering customers at the industrial and enterprise level to utilize maker technology in their advanced projects.

News ID 4914

Hall-Stand 4882

RTI to present Connext DDS middleware for smart distributed applications

At embedded world 2017, Real-Time Innovations will present its connectivity platform for real-time applications in the Industrial Internet of Things. The focus will be on the RTI Connext DDS middleware for smart distributed applications, its application in the field of autonomy as well as the security specification Connext DDS Secure. Connext DDS live demos can also be seen at partner booths. RTI Connext DDS Secure provides the world's first off-the-shelf messaging platform for the IIoT, which also complies with the new DDS (Data Distribution Service) Security specification from the Object Management Group.

News ID 4882

Hall-Stand 4955

Rohde & Schwarz: powerful T&M equipment at embedded world 2017

Exceptional quality and innovation doesn't have to be expensive. And Rohde & Schwarz proves this at embedded world in Nuremberg with its new entry-level T&M solutions. One of the highlights at the Rohde & Schwarz booth at embedded world 2017 is the new 6 GHz model of its extremely powerful R&S RTO2000 lab oscilloscopes. Their ability to handle multidomain applications makes the compact lab oscilloscopes perfect for IoT applications. Their outstanding characteristics are ideal both for demanding T&M tasks, such as debugging fast communications interfaces, and for the increasingly important power integrity measurements.

Rohde & Schwarz has also introduced two new probes specifically designed for power integrity measurements for use with its R&S RTE and R&S RTO2000 oscilloscopes. The R&S RT-ZPR20 is a highly sensitive and extremely low-noise power rail probe with a bandwidth of 2 GHz. The large offset range of ±60 V enables analysis of the smallest signal interferers during power integrity measurements, even on DC power supplies with a high voltage level. The R&S RT-ZP1X passive 1:1 probe is especially suitable for voltage measurements in the millivolt range.

Rohde & Schwarz continually expands the functionality offered by its R&S Scope Rider, which now combines eight measurement instruments into one robust, portable oscilloscope, making it the most versatile oscilloscope in its class. New features include a spectrum analyzer, a frequency counter and a harmonic analyzer for evaluating the quality of power supplies. A new CAN/LIN trigger and decoder option is also available for automotive applications. Users always have the tools they need, both on the road when performing installation, service and maintenance activities and in the lab. Since the fully insulated, handheld instrument meets measurement category CAT IV requirements, it can be used to perform measurements at the source of low-voltage installations up to 600 V.

News ID 4955

Hall-Stand 4931

Rohde & Schwarz: 6 GHz lab oscilloscope for multi-domain applications

The new R&S RTO2000 model with 6 GHz bandwidth allows developers to test the radio interfaces of 802.11ac WLAN components for IoT modules in the 5 GHz band as well as fast communications interfaces such as USB 3.1 Gen 1 with data rates of 5 Gbit/s. Thanks to its multi-domain functionality, only a single compact instrument is needed to analyze power supplies, processors and sensors with up to 6 GHz bandwidth. Synchronized time, frequency, protocol and logic analyses results allow users to debug at the system level.

News ID 4931

Hall-Stand 4877

Cadence announces availability of Bluetooth 5 verification IP

Cadence Design Systems announced the immediate availability of its Cadence Verification IP (VIP) for Bluetooth 5, the industry's first VIP for the latest version of Bluetooth technology. Bluetooth 5 increases data broadcasting capacity by 800 percent, quadruples the range and doubles the connection speed of low-energy devices to deliver seamless, short-range mobile connectivity. Cadence also strengthened its IP portfolio with the addition of VIP for Bluetooth 4.2 and the Cadence TripleCheck productivity tool for Bluetooth 5.

News ID 4877

Hall-Stand 4862

Green Hills: high assurance platforms for automotive cockpit with Renesas' R-Car

Green Hills Software has announced an expanded collaboration with Renesas Electronics by delivering integrated automotive solutions and global customer adoption of the safe and secure INTEGRITY real-time operating system and INTEGRITY Multivisor virtualization for the latest Renesas R-Car automotive computing system-on-chips. Both companies have demonstrated solutions based on Renesas' R-Car H3 and R-Car D1 SoCs and new support for R-Car M3, showcasing solutions for advanced driver assist systems, reconfigurable digital instrument clusters, eCockpit domain consolidation and connected car V2X.

News ID 4862

Hall-Stand 4913

Green Hills: DLM simplifies secure device manufacturing with managed service

INTEGRITY Security Services has announced the Device Lifecycle Management (DLM) Managed Service to simplify manufacturing cryptographic IoT devices requiring digitally signed software, zero-exposure key generation, certificate authority, and automotive vehicle-to-vehicle certificate generation. The DLM Managed Service eliminates the overhead of developing and supporting internal PKI systems for secure delivery of digital trust assets to partners and suppliers across global supply chains.

News ID 4913

Hall-Stand 4939

Wibu-Systems: CodeMeter Embedded 2.0 makes its debut at Embedded World 2017

The next generation of CodeMeter Embedded will be introduced by Wibu-Systems at Embedded World. CodeMeter Embedded 2.0 is a runtime environment that supports the complete array of platforms and operating systems intelligent device manufacturers rely upon, including Intel x86, PPC, ARM, Windows, Linux, Android, VxWorks, and QNX. The modular and scalable design enables developers to integrate only the functionality they need and create a software solution with the smallest footprint possible, the highest level of security and the most versatile licensing, all without affecting performance.

News ID 4939

Hall-Stand 4924

Rigol: new function generators and decode options for oscilloscopes

At embedded world, Rigol Technologies will present innovations and proven instruments for test and measurement applications. With the new DG1022Z model, a part of the new family of fast, easy-to-use, and ultra-modern test devices, Rigol presents a multifunctional generator that enables a wide range of tests by combining multiple functions in one instrument. The devices combine application fields like function generator, arbitrary waveform generator, pulse generator, harmonic generator, and analog/digital modulation as well as a counter function. In addition to the Direct Digital Synthesizer technology, which provides very stable, precise, and slightly distorted signals, the signal jitter is reduced to 200ps by SiFi, the new and innovative Signal-Fidelity technology. The device is characterized by 25 MHz bandwidth for sine and square signals and two standard, fully functional and independently operating channels. Up to 160 integrated waveforms are available and the various modulation modes include AM, FM, PM, ASK, FSK, PSK and PWM. The standard interfaces USB Host & Device, and LAN (LXI CoreDevice 2011) are available for communication purposes.

News ID 4924

Hall-Stand 4956

Renesas: next-gen emulator contributes to reduced Embedded software development time

Renesas Electronics announced the E2 Emulator, a new-generation on-chip debugging emulator. The E2 is intended as a development environment for the latest devices in the Renesas RH850, RX, and RL78 Families of microcontrollers, and for a selection of automotive SoCs. The new emulator supports the extended debugging functionality of the RH850 Family and contributes to shortening the time required for CAN communication debugging and the current consumption debugging.

News ID 4956

Hall-Stand 4902

Axiomtek: micro-sized DIN-rail IIoT gateway for entry level applications

Axiomtek introduces the ICO100, a DIN-rail fanless embedded system designed to support the low-power Intel Atom x5-E3930 processor. This DIN-rail industrial IoT gateway supports one DDR3L SO-DIMM slot with up to 8 GB system memory. Its proven rugged construction ensures reliable operation in harsh environments with an extended temperature range of -20 to 70°C and anti-vibration up to 2G.

News ID 4902

Hall-Stand 4943

MicroSys: operate machines in real time, locally and web based

XiSys has created a so-called Hybrid Graphics Server, that offers Web based Man Machine Interaction remotely on smart devices or systems running a browser and local human interaction directly in front of a machine on the control device. The RTOS Microware OS-9 in its newest version 6.1 complements those capabilities perfectly as it is designed for applications relying on reliable operation, precision, fast boot up sequences and shortest interrupt latencies. The architecture of OS-9 is clearly structured and allows slim run time environments to address cost effective system designs. The XiBase9 & RTOS Microware OS-9 6.1 combines WEB visualization via standard browsers, like Edge, Chrome, Firefox, Internet Explorer, Opera or Safari; recourse to HTML5 based standards like web sockets, canvas, objects and java script do not require any system overhead on device level; visualization of graphic objects within a browser is performed without any apps, plug ins, cookies, java programs and by the support of SH1 encryption secure operation can be achieved.

News ID 4943

Hall-Stand 4876

Rutronik: everything for Embedded IoT applications

Rutronik Elektronische Bauelemente is focusing on the Internet of Things at embedded world 2017. RUTRONIK EMBEDDED offers visitors selected wireless chips and mod-

More information about each news is available on www.Embedded-Control-Europe.com/magazine You just have to type in the "News ID". —

boards & solutions ECE

Editors
Jürgen Hübner
phone +49(0)8092-2477413
fax +49(0)8092-2477429
jh@iccmedia.com

Wolfgang Patelay
wp@iccmedia.com

Tony Devereux
devrex@teyboyz.freeserve.co.uk

For Reader Inquiries and Address Changes **please contact:**
info@iccmedia.com

Sales & Marketing Director
Manfred Blumoser
phone +49(0)8092-2477411
fax +49(0)8092-2477429
mb@iccmedia.com

Claudia Mellein
cm@iccmedia.com

Christiane Lockner
cl@iccmedia.com

Sales Office - UK and USA, Benelux, Scandinavia
Blue Sky Communications
Malcolm Cameron
21 Cliffe Avenue
Westbrook,
Margate, Kent CT9 5DU, UK
phone +44 (0)77 88-10 84 11
fax +44 (0)80 82-8010 57
mc@blue-sky-communications.com

Sales Office - Asia
Jean Cheng, *jean@i-media.com.tw*
Vivian Hung, *vivian@i-media.com.tw*
Innovative Media Information & Services
7F-3, No. 26, Sec. 2, Ming-Quan East Rd.
Taipei 104 Taiwan
phone +886 2 2563 1186

Head Office

ICC Media

ICC Media GmbH
Rauwagnerstr. 5
85560 Ebersberg / Germany

Editorial Office UK
36a Blackacre Road
Theydon Bois
Essex, CM16 7LU

ules, displays, boards, and storage solutions characterized by their durability, long-term availability, and high integration capacity for industrial and IoT applications. While RUTRONIK SMART features radio technology, sensors, microcontrollers, power management, battery management, and cryptography ICs that are optimized for internet protocol usage, are smaller in size, have a lower power consumption, and offer a high degree of integration. In cooperation with Chemnitz University of Technology and the Chair for Measurement and Sensor Technology (MST), Rutronik presents the latest research results from the field of impedance measurement based on the impedance spectroscopy of batteries. Researchers from the University will be on hand throughout the exhibition to answer any questions you may have. Booth highlights include smart Embedded battery management, sensing, visualization, data processing, storage, the Intel Joule Developer Kit and communication.

Experienced product managers and field application engineers from all specialist areas will be available to provide competent support and advice. In addition, Rutronik will present Rutronik24 as a distributor for SMEs and large businesses with mid-range component requirements. Selected manufacturers will showcase their current top products and be happy to answer any questions concerning products and applications.

News ID 4876

Hall-Stand 4930

NI partners with IBM and SparkCognition to advance the IIoT

NI announced that SparkCognition is partnering with NI and IBM to collaborate on the Condition Monitoring and Predictive Maintenance Testbed. The goal of the collaboration is to deliver an unprecedented level of interoperability among operational technology and informational technology as organizations search for better methods to manage and extend the life of aging assets in heavy machinery, power generation, process manufacturing and a variety of other industrial sectors.

News ID 4930

Hall-Stand 4861

Advantech joins LoRa Alliance to drive Industry 4.0 and smart city applications

Advantech has joined the LoRa Alliance, which is committed to developing a new industrial standard to empower the Internet of Things. Advantech will provide the LoRa gateway and sensor node devices based on the M2.COM open standard for Industry 4.0 and Smart City applications.

News ID 4861

Hall-Stand 4936

Logic Technology: integrating rich GUI into your next IOT application

NXP invites you to join in on their webinars on 19.02.2017, where TouchGFX will be one of three parts. Based on the high-performance ARM Cortex-M4 core, the LPC54000 series of single-core and dual-core MCUs represents the next-generation of industry-leading power efficiency. With added scalability and feature integration, these breakthrough devices enable continued market growth in today's competitive environment.

News ID 4936

Hall-Stand 4927

IntervalZero plans upgrade-path for WinCE users to Windows 10 IoT

IntervalZero announces its latest Real Time solution for Windows 10 IoT. WinCE is often used for headless products and in high volume applications. With the move of Microsoft to Windows 10 only 3 versions are available. Win10 IoT Enterprise, Win10 IoT Mobile and Win10 IoT Core. The last version could be an upgrade path for WinCE-users, used to very low license costs. But if there is an exact timing or determinism necessary, this possibility lacks the real time part.

News ID 4927

Hall-Stand 4925

ADLINK introduces high performance machine vision products

ADLINK Technology announced the release of its new NEON-1021 Intel Atom E3845 processor-based smart camera, EOS-1300 4CH PoE Compact Vision System, and PCIe-GIE72/74 2/4CH GigE Vision PoE+ frame grabbers, delivering precise image acquisition with high integration, effectively reducing TCO.

News ID 4925

Hall-Stand 4954

Nexperia emerges as dynamic new force in discretes, logic and MOSFETs

Nexperia announced the formal completion of its launch as a separate entity. Headquartered in Nijmegen, Netherlands, Nexperia is a stand-alone leader in discretes, logic and MOSFETs, retaining all the expertise, manufacturing resources and key personnel of the former NXP division, while bringing a new focus and powerful commitment to these product areas.

News ID 4954

Advertisers Index

ADLINK
Leading EDGE COMPUTING

Low Power Computer-on-Modules

» Die neuesten Intel® SoCs auf SMARC Modulen

» Technischer Support in Ihrer Region

LEC-AL

SMARC® Short Size Modul mit 6. Gen. Intel® Pentium™ und Intel® Celeron® Prozessoren

- 3x Display Anslüsse: Dual channel LVDS, HDMI/DP++, DP++
- 2x MIPI CSI Kamera
- 1x GbE mit IEEE1588 Unterstützung
- 1x SATA 3.0, eMMC
- Erweiterter Temperaturbereich: -40°C bis +85°C
- SMARC 2.0 kompatibel

LEC-BW

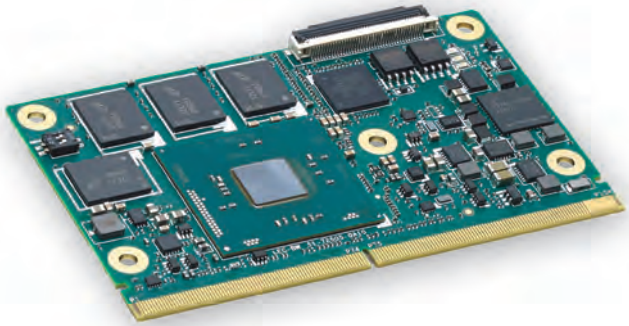

SMARC® Short Size Modul mit Intel® Pentium™ und Celeron® N3000 Series SoC

- LVDS, HDMI
- 2x MIPI CSI camera interfaces, GbE, 3x PCIe
- 1x SDIO, 2x SATA3 6 Gbit/s, 1x eMMC
- 1x USB 3.0, 3x USB 2.0, 12x GPIO, 2x SPI, 4x I²C
- SMARC 1.1 kompatibel

ADLINK TECHNOLOGY GmbH

- *Tel: +49-621-43214-0*
- *EMEA@adlinktech.com*
- *www.adlinktech.com*

intel IoT Solutions Alliance Premier