Embedded COMPUTING DESIGN

JUNE 2016
VOLUME 14 | 4
EMBEDDED-COMPUTING.COM

EMBEDDED LENS PG 6
A holistic approach to IoT

IOT INSIDER PG 8
The software and sensors of Amazon Echo

Top Embedded Innovators PG 32

2016 TOP INNOVATIVE PRODUCT NOMINEES Embedded COMPUTING DESIGN PG 38

PLUS:
HOW 3D PRINTING WILL ROCK THE WORLD
PG 24

JEFF BADER
Micron Technology, Inc.

KENT MEYER
Emcraft Systems

TYSON TUTTLE
Silicon Labs

CONTENTS

Embedded COMPUTING DESIGN

June 2016 | Volume 14 | Number 4

opsy.st/ECDLinkedIn @embedded_comp

FEATURES

10 **Delivering more natural, personalized, and secure voice control for today's connected world**
By Rammy Bahalul, VocalZoom

14 **FPGAs and audio processors enable unique industrial applications**
By Ted Marena, Microsemi

18 **Mixing C and Java programming in embedded, IoT designs**
By Vincent Perrier, MicroEJ

22 **Feel "free" to use embedded Linux**
By Chris Simmonds, Looking after the Inner Penguin

24 **3D printing explained, and how it will rock the world**
By John Hornick, www.3DPrintingWillRockTheWorld.com

28 **Open source opens many licensing issues for 3D printing**
By Maya M. Eckstein and Eric J. Hanson, Hunton & Williams LLP

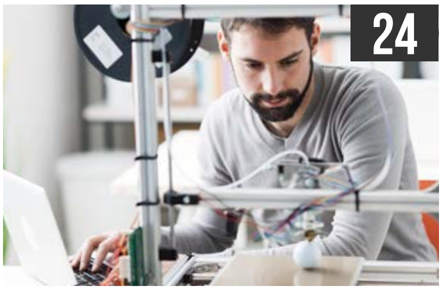

2016 TOP EMBEDDED INNOVATORS

32 **Tyson Tuttle, CEO, Silicon Labs**

34 **Kent Meyer, CEO, Emcraft Systems**

36 **Jeff Bader, Vice President, Micron Technology, Inc.**

38 TOP INNOVATIVE PRODUCT NOMINEES

COLUMNS

5 **TRACKING TRENDS**
The beginnings of innovation in the "new" embedded
By Rory Dear, Technical Contributor

6 **EMBEDDED LENS**
A holistic approach to IoT
By Curt Schwaderer, Editorial Director

8 **IOT INSIDER**
Deconstructing Alexa – Software and sensors of the Amazon Echo and beyond
By Brandon Lewis, Technology Editor

43 **EDITOR'S CHOICE**
By Curt Schwaderer, Editorial Director

COVER

From integrated silicon to educational reform, the 2016 Top Embedded Innovators section features industry standouts Jeff Bader of Micron Technology, Kent Meyer of Emcraft Systems, and Tyson Tuttle of Silicon Labs.

WEB EXTRAS

> **Predixions and precision for the Industrial Internet**
By Brandon Lewis, Technology Editor
http://bit.ly/IES2016FT

> **Let crowdfunding help fund your project**
By Jeremy Cook, Engineering Consultant
http://bit.ly/JCCrowdfunding

> **Don't let hackers in the JTAG port**
By Mike Rohrmoser, Digi International
http://bit.ly/JTAGHacker

> **Security is in the slow lane when it comes to connected cars**
By Lancen LaChance, GlobalSign
http://opsy.st/28XrsNu

DOWNLOAD THE APP

Download the Embedded Computing Design app

iTunes: itun.es/iS67MQ
Magzter: opsy.st/ecd-magzter

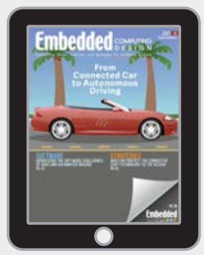

Published by:

OpenSystems media.

ISSN: Print 1542-6408
Online: 1542-6459

AD LIST

PAGE	ADVERTISER
31	**ACCES I/O Products, Inc.** – Saving space, the final frontier
9	**Advantech Embedded Group** – Design-in embedded solutions
44	**American Portwell Technology** – Empowering the connected world
13	**BZ Media** – InterDrone: The international drone conference and exposition
17	**COMMELL Systems Corporation** – Leading provider of SBCs
16	**Interface Concept** – Rugged HPEC boards for your OpenVPX systems
21	**Technologic Systems** – Don't let a natural disaster become a total disaster
3	**WinSystems, Inc.** – Delivering critical data connectivity

SOCIAL

Facebook.com/
Embedded.Computing.Design

@Embedded_comp

LinkedIn.com/in/
EmbeddedComputing

Pinterest.com/Embedded_Design/

Instagram.com/
Embedded_Computing

youtube.com/user/
VideoOpenSystems

Embedded COMPUTING DESIGN

EMBEDDED COMPUTING BRAND DIRECTOR Rich Nass rnass@opensystemsmedia.com
EMBEDDED COMPUTING EDITORIAL DIRECTOR Curt Schwaderer cschwaderer@opensystemsmedia.com
TECHNOLOGY EDITOR Brandon Lewis blewis@opensystemsmedia.com
TECHNICAL CONTRIBUTOR Rory Dear rdear@opensystemsmedia.com
CONTENT ASSISTANT Jamie Leland jleland@opensystemsmedia.com
DIRECTOR OF E-CAST LEAD GENERATION AND AUDIENCE ENGAGEMENT Joy Gilmore jgilmore@opensystemsmedia.com
CREATIVE DIRECTOR Steph Sweet ssweet@opensystemsmedia.com
SENIOR WEB DEVELOPER Konrad Witte kwitte@opensystemsmedia.com
WEB DEVELOPER Paul Nelson pnelson@opensystemsmedia.com
DIGITAL MEDIA MANAGER Rachel Wallace rwallace@opensystemsmedia.com
CONTRIBUTING DESIGNER Joann Toth jtoth@opensystemsmedia.com

SALES

SALES MANAGER Tom Varcie tvarcie@opensystemsmedia.com (586) 415-6500
STRATEGIC ACCOUNT MANAGER Rebecca Barker rbarker@opensystemsmedia.com (281) 724-8021
STRATEGIC ACCOUNT MANAGER Bill Barron bbarron@opensystemsmedia.com (516) 376-9838
STRATEGIC ACCOUNT MANAGER Eric Henry ehenry@opensystemsmedia.com (541) 760-5361
STRATEGIC ACCOUNT MANAGER Kathleen Wackowski kwackowski@opensystemsmedia.com (978) 888-7367
SOUTHERN CALIFORNIA REGIONAL SALES MANAGER Len Pettek lpettek@opensystemsmedia.com (805) 231-9582
SOUTHWEST REGIONAL SALES MANAGER Barbara Quinlan bquinlan@opensystemsmedia.com (480) 236-8818
NORTHERN CALIFORNIA REGIONAL SALES MANAGER Twyla Sulesky tsulesky@opensystemsmedia.com (408) 779-0005
ASIA-PACIFIC SALES ACCOUNT MANAGER Elvi Lee elvi@aceforum.com.tw
EUROPE SALES ACCOUNT MANAGER James Rhoades-Brown james.rhoadesbrown@husonmedia.com

OpenSystems media ®
WWW.OPENSYSTEMSMEDIA.COM

PUBLISHER Patrick Hopper phopper@opensystemsmedia.com
PRESIDENT Rosemary Kristoff rkristoff@opensystemsmedia.com
EXECUTIVE VICE PRESIDENT John McHale jmchale@opensystemsmedia.com
EXECUTIVE VICE PRESIDENT Rich Nass rnass@opensystemsmedia.com
CHIEF TECHNICAL OFFICER Wayne Kristoff
GROUP EDITORIAL DIRECTOR John McHale jmchale@opensystemsmedia.com
VITA EDITORIAL DIRECTOR Jerry Gipper jgipper@opensystemsmedia.com
PICMG EDITORIAL DIRECTOR Joe Pavlat jpavlat@opensystemsmedia.com
MANAGING EDITOR Jennifer Hesse jhesse@opensystemsmedia.com
ASSISTANT MANAGING EDITOR Lisa Daigle ldaigle@opensystemsmedia.com
SENIOR EDITOR Sally Cole scole@opensystemsmedia.com
ASSOCIATE EDITOR Mariana Iriarte miriarte@opensystemsmedia.com
CREATIVE PROJECTS Chris Rassiccia crassiccia@opensystemsmedia.com
FINANCIAL ASSISTANT Emily Verhoeks everhoeks@opensystemsmedia.com
SUBSCRIPTION MANAGER subscriptions@opensystemsmedia.com

CORPORATE OFFICE
16626 E. Avenue of the Fountains, Ste. 201 • Fountain Hills, AZ 85268 • Tel: (480) 967-5581

SALES AND MARKETING OFFICE
30233 Jefferson • St. Clair Shores, MI 48082

REPRINTS

WRIGHT'S MEDIA REPRINT COORDINATOR Wyndell Hamilton whamilton@wrightsmedia.com (281) 419-5725

rdear@opensystemsmedia.com

THE BEGINNINGS OF INNOVATION IN THE "NEW" EMBEDDED

By Rory Dear, Technical Contributor

I'm often guilty of tunnel vision when it comes to embedded, as the majority of my career has been spent where embedded and industrial were synonymous, and often interchangeable. However, events like Computex 2016 always remind me just how wide the scope of embedded is today. In fact, I now genuinely believe that our industry can no longer be described as niche – we've finally become mainstream.

With that brings an abundance of new opportunities for all involved, though new prey attracts new predators. The enterprise computing behemoths are seizing the opportunity to jump on the embedded bandwagon where arguably traditional vendors have been slow to present real solutions for the opportunities created by the revolution of the Internet of Things (IoT) and its umbrella terms. Traditional embedded vendors have been pushed to rapidly diversify their offerings and prove to customers they truly understand the connected future; the days of slapping IoT or Industry 4.0 across one's booth without clear evidence of that fact are gone.

This diversification is inspiring, where exciting new use cases are emerging daily by opening embedded to a wider inexperienced audience. And there's also a fair amount of dross to sift through. I saw one of these new IoT devices at the venue for startups at Computex, InnoVEX – the Pet Cam. It simplistically allows one to remotely monitor their pet via an IP camera. Nothing new there, but goes much further by pushing notifications of their detected emotional state to your smartphone! One can even appease an (allegedly) negative state by dispensing treats remotely or activating a light-chase game.

Before arriving at InnoVEX, I read an amusing statement that I was keen to disprove. "IoT is like teenage sex: Everyone talks about it, nobody really knows how to do it. Everyone thinks everyone else is doing it, so everyone claims they are doing it too." My specific interest was to compare the primary Computex hall of long established vendors to the fresh blood at InnoVEX, and I was pleasantly surprised at both.

To me the startups peddling their wares at InnoVEX shared true vision, invariably created by young and enthusiastic engineers without preconceptions – a world where anything is possible with technology. Innovation desperately needs the fresh thinking of youth and they inherently understand the concept. But I fear that realization of their end-to-end concepts isn't achievable on their own. That said, as I walked around the InnoVEX exhibition, it felt more like a marketplace to advertise these start-ups to prospective buyers who have the clout to truly realize these young dreams.

FINALLY, IT SEEMS THAT VENDORS HAVE THOUGHT LONG AND HARD ABOUT WHERE THEIR POSITION IS IN THIS NEW MAINSTREAM EMBEDDED MARKET.

At Computex proper I felt vendors had realized that cementing their place as the market leader in a specific embedded niche makes far better business sense than purporting to carry a perceivably weak end-to-end solution, particularly around IoT. Such solutions are only realistically achievable by the industry behemoths, though some of these concepts are new to them.

Finally, it seems that vendors have thought long and hard about where their position is in this new mainstream embedded market. So hopefully the postulating and jostling are over and innovation will start to snowball.

cschwaderer@opensystemsmedia.com

A HOLISTIC APPROACH TO IoT

By Curt Schwaderer, Editorial Director

The Internet of Things (IoT) has caught fire and many industries are seeing initial rollouts of these solutions that promise to provide significant value to both customer and vendor. The natural reaction is to treat this as a technology initiative, focusing on how best to bring together data from smart, connected devices. And in certain cases, this approach has merit.

Bsquare has a somewhat unique view on IoT – focusing on business value and taking a holistic approach to integrating IoT into existing processes and workflows. The result is a faster rollout, quicker feedback loop, and meaningful business improvements that lower risk associated with new technology adoption.

Holistic IoT approach

Bsquare has been around for 22 years with a traditional core business revolving around embedded integration and deployment of solutions that implement business objectives for Fortune 500 companies. This 22-year history involves all the usual suspects – embedded boards, RTOSs, and BSPs, along with embedded control software to run them. Over time these devices have all become connected, and Dave McCarthy, Senior Director of Products at Bsquare, provided an example of the holistic IoT approach.

"One good example of a holistic IoT approach was the Coca Cola Freestyle machine," McCarthy said. "132 flavors with a touchscreen user interface that allows you to mix and match your flavors before you fill your glass. This project proceeded like any other networked embedded system solution, but this one went further. Data generated from the machine was being used to understand consumer behavior and optimize the supply chain."

McCarthy also mentioned similar projects involving Costa coffee machines – barista quality coffee in a vending machine where Bsquare did the software work, from board to applications. "All of these cases were on a consulting and services basis and we learned that the motivations for IoT extended beyond just extracting and transmitting device data."

Common functions and use case focus

McCarthy mentioned that one of the values of starting from the consulting space was the emphasis on business use cases. This proved invaluable when Bsquare began to turn attention to an IoT framework and associated system requirements, of which there are three (Figure 1):

1. **IoT device enablement.** Sensors connected to local processing provides data as well as intelligence at the device level, which is key for distributing event information upstream while also accepting a feedback loop of actions in order to respond to business requirements and use cases. Sometimes it's ok to push data to the cloud, but in many cases this isn't practical as slow or no connectivity can mean excess network transmission costs associated or additional latency that could be a deal breaker for a real-time control system, for example.
2. **Real-time monitoring and event processing.** This encompasses enterprise and public data sources. With the disparate nature of IoT information, a powerful real-time monitoring and event processing environment enables the most efficient use of all available data sources.
3. **Analytics.** The connections between analytics, automation, and the device can be centralized in an on-premise server or cloud environment, and enables the simplification and distribution of IoT data.

The DataV stack

DataV is a holistic solution that includes all of the aforementioned components with solid architectural separation between the device and data center elements (Figure 2). However, the real-time and data analytics components are architected to be able to be flexibly moved between on-premises and remote cloud environments.

The distributed logic between device and cloud can be divided into a baseline OEM layer that enables customers to write their own rule sets. This gives IoT device manufacturers the ability to differentiate through a flexible architecture that puts information processing where it's most effective – on the device, on-premise, or in the data center.

The third layer of the stack provides optimized analytics and resulting predictions that feed back into the lower layers for incremental value and optimized intelligence.

IoT is about actionable intelligence

The DataV architecture is about creating a closed loop with actionable intelligence on one end and business objectives on the other. For example, a logistics use case demands constant uptime for delivery trucks, as a truck out of service can't deliver cargo. Therefore, predictive, rather than reactive, maintenance is preferred.

A Bsquare client already had a solution for collecting telematics data from its trucks, so Bsquare augmented that information with surrounding operational data (such as repair history) with the goal of performing automatic root cause analysis. To drive root cause analysis execution, the solution considered reported error codes, potential fixes based on those codes, as well as additional historical and real-time data that could be used to eliminate a subset of the proposed fixes. Once that process was complete, a probability analysis was performed on the remaining potential fixes, and those findings were then combined with part inventory and technician skill level to dramatically cut diagnostics time, improve the quality of the fix, and lower time to repair. DataV was central to the success of this "aftermarket" IoT deployment.

IoT and retrofits

IoT is fine for new systems, but what about the massive amount of legacy equipment? Smart vending is one such example that McCarthy described.

"There is a massive amount of fielded equipment not ready for retirement. But there is high value for participation in a connected world," McCarthy said. "Bsquare is part of the Intel Internet of Things Solutions Alliance, so we worked with Intel on a hardware solution that can bridge the gap. The board uses Linux and DataV to bridge the basic metrics provided by legacy vending equipment and adds additional sensor information with local monitoring and automation capabilities. This is a great example of enabling IoT without needing to 'rip-and-replace' existing equipment."

Summary

IoT is about customer experience and the ability to use technology to monitor, analyze, and automate in order to provide new and better services to those customers. Starting with defined business rules and integrating the technology to move toward improving service can yield tangible results in a way that clearly justifies the investment.

While many solutions are designed to solve a piece of the problem, a holistic approach takes into account all of the necessary components to achieve business use cases so that value can be realized more quickly and with less risk.

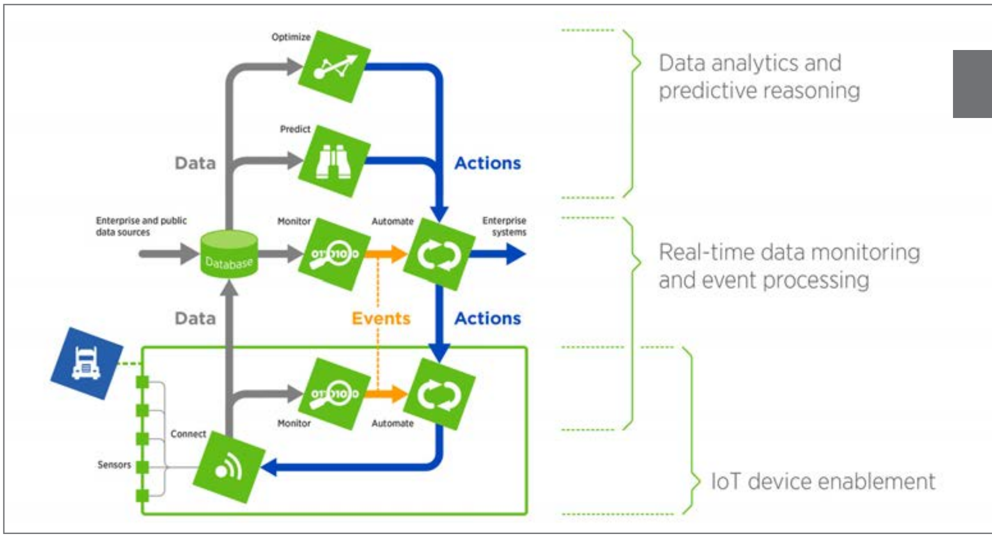

FIGURE 1

The three requirements of a successful IoT system.

FIGURE 2

The DataV architectural stack.

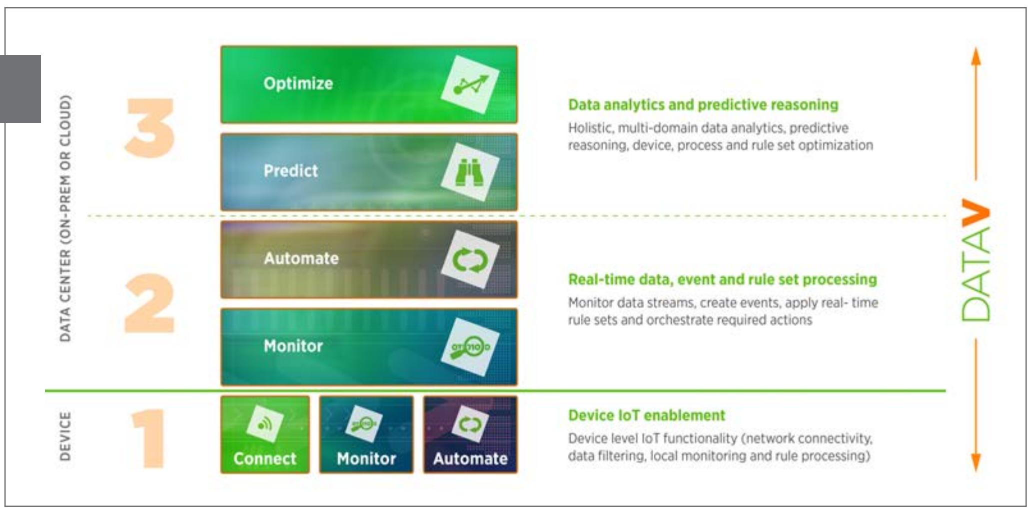

DECONSTRUCTING ALEXA – SOFTWARE AND SENSORS OF THE AMAZON ECHO AND BEYOND

By Brandon Lewis, Technology Editor

The Amazon Echo is the epitome of an Internet of Things (IoT) device. It combines an embedded applications processor from Texas Instruments, MEMS microphones from Knowles, Wi-Fi and Bluetooth wireless connectivity, an AWS cloud backend, and support for diverse applications. It's also multi-function, which increases the platform's value for consumers (bundled services), as well as Amazon (multi-dimensional insights into customer behavior and trends). The glue that ties all of this together is, of course, software.

The Echo's signature feature, automatic speech recognition (ASR), is enabled by software algorithms that not only provide the language modeling and natural language understanding capabilities that make the platform unique, but also help offset the rigors of reverberant speech. Reverberant speech is a phenomenon that occurs in indoor environments when an audible signal reflects or bounces off of various surfaces, creating noise in the form of echoes that diminish the direct path signal from speaker to microphone. As you can imagine, this wreaks havoc on speech recognition, but consider the real-world use case of the Amazon Echo wherein reverberant speech is often the *only* signal available from a speaker communicating with the device.

Jeff Adams, CEO of Cobalt Speech & Language, Inc. and former Senior Manager of the speech and language groups at Amazon, worked on the Echo. He attributes the platform's success in situations where his wife yells, "'Alexa, what time is it?' and hears the answer even though she's three rooms away, down the hall, and around the corner" to cloud-based deep neural networks (DNNs) capable of performing roughly 1 billion arithmetic operations per second in support of ASR algorithms, beamforming, and noise cancellation techniques. But, while Adams suggests that kind of computing power became possible after cycles of Moore's law and could at some point be available on processors beyond the data center, those performance requirements don't leave much hope for accurate ASR today in embedded devices not backed by the power of the cloud.

Sensors, software, and embedded speech rec

Even though acoustic and language processing models such as those used for the Echo can be compressed, the reality is that compression comes with tradeoffs. The more ASR models are compressed the less accurate they become, and typically the size of language libraries shrinks dramatically from the linguistic openness of platforms like the Echo to perhaps a few hundred or a few thousand words. Furthermore, even after compression you're probably still talking about hundreds of MB for such models, which is a huge burden on even high-end smartphones.

However, innovations in sensor technology are emerging that could help remove some of the overhead associated with massive DNNs, namely the use of multiple, heterogeneous inputs. For instance, Cobalt is partnering with human-to-machine communications (HMC) company VocalZoom, a manufacturer of optical sensors that pair with acoustic microphones to eliminate background noise and improve directional acquisition for speaker isolation.

The optical sensor technology works by converting vibrations from a speaker's cheek, larynx, and other facial areas into an audio signal, though one devoid of background noise due to the low frequencies at which skin vibrates. This information is then fused with inputs from traditional acoustic microphones to generate noise-free audio signals that can be leveraged in the absence of cloud-based DNNs to reduce the effects of reverberant speech, and even enable applications such as access control and voice authentication. For example, such an implementation could prevent systems like the Echo from waking up when a TV commercial mentions "Alexa" (more on optical sensors can be found in "Delivering more natural, personalized, and secure voice control for today's connected world" on page 10).

Additionally, Adams says that other sensors are starting to be considered in the ASR equation, particularly as his company works towards speech *classification* engines designed to infer background information about a speaker, such as age, gender, physical and emotional state, and even possibly to aid in early diagnosis of medical conditions like Parkinson's and Alzheimer's. Cameras and inputs from medical devices would be obvious complements in these types of applications, which could lead to the next level of sensor data fusion for the Internet of Things.

DELIVERING MORE NATURAL, PERSONALIZED, AND SECURE VOICE CONTROL FOR TODAY'S CONNECTED WORLD

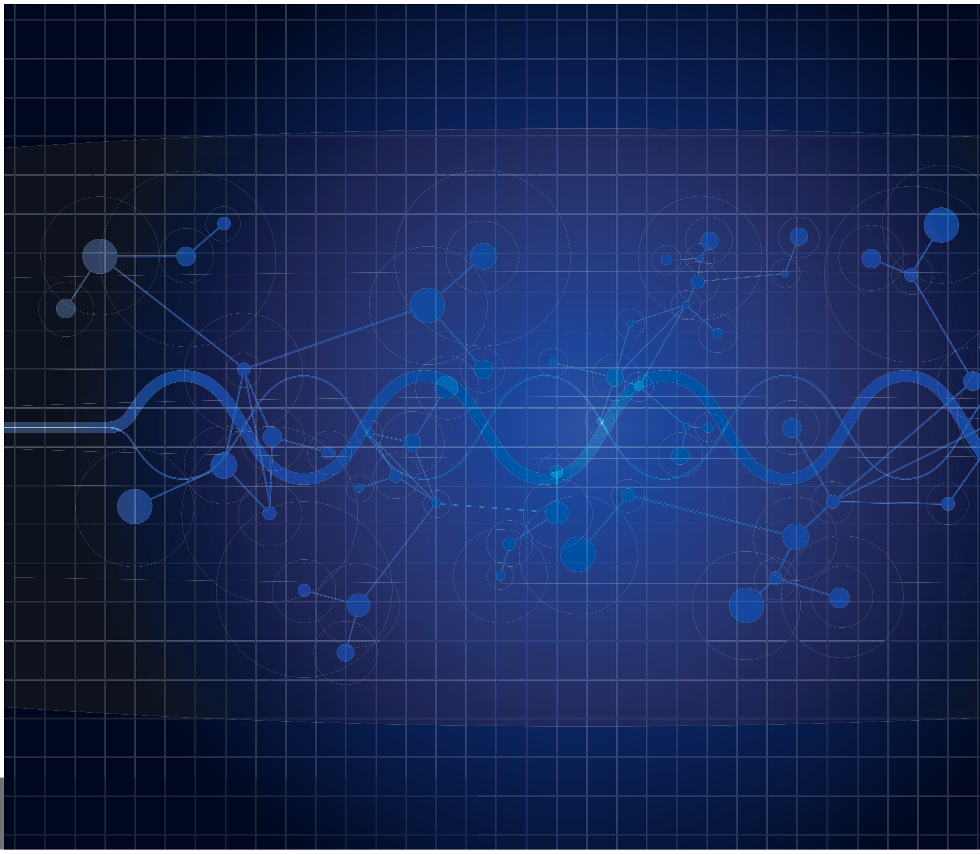

By Rammy Bahalul

The industry has moved from punch card to keyboard and from mouse to touchscreen, all in pursuit of more direct system manipulation to optimize user experience (UX) in today's increasingly mobile and interconnected world. These are all abstractions of physical devices, though, and voice control has been heralded as the next step toward a more natural UX. Unfortunately, today's solutions can't deliver what machines need to understand – people – resulting in poor performance and no convenient way to control a new generation of voice-only products and services.

One of the biggest impediments to satisfactory voice control performance has been ambient noise, including nearby conversations, outdoor sounds, and reverberation when speaking in certain indoor environments. The use of multiple acoustic microphones and microphone arrays to improve directional acquisition has proven expensive and incapable of adequately isolating the speaker for reliable voice control. Now, a new approach is available that leverages optical lasers and interferometry techniques to gather additional critical information exclusively about the user communicating with a device. Combining this optical information with the output from an acoustic microphone gives automatic speech recognition (ASR) engines something they have never had before – a near-perfect reference audio signal directly from the speaker's facial vibrations, regardless of noise levels.

Understanding the special challenges of human-to-machine communications

Human-to-machine communications (HMC) technology enables humans to interact with and control a variety of networked devices as quickly and efficiently as possible. While voice is an excellent interface, the problem with using today's ASR engines for HMC applications is that they have generally been designed for human listeners and typically only perform well if words are spoken clearly and there is no background noise. This, of course, is not the case in a real-world, noisy environment. Machines are incapable of inferring meaning as humans do if background noise periodically drowns out the speaker, and while voice-recognition software can be trained to understand accents and other speech patterns, they cannot be trained to ignore background noise. Solutions must be able to isolate the speaker's voice from others in the background, as well as from other types of ambient noise.

In tests of voice-recognition solutions in a moving vehicle with windows fully open and with speakers in the background, the word/command recognition rate typically drops to 0 percent (Figure 1). The industry has pursued a number of approaches to solving this problem over the past 20 years but, in general, these efforts have delivered only single-digit percentage improvements in word recognition performance.

Reducing or eliminating background noise to isolate the speaker's voice is critical to improving the accuracy of automatic speech recognition engines in HMC applications. Acoustic microphone technology alone does not provide enough directional acquisition capability to achieve this level of speaker isolation, even with multiple microphones and microphone arrays. However, if the output from an acoustic microphone can be paired with additional outputs associated exclusively with the speaker, there is an opportunity to reduce word error rates by at least 60 percent.

Applying optical laser technology and interferometry techniques

The key to improving voice recognition using optical laser technology is the ability to measure the distance and velocity of facial vibrations during speech. This approach takes advantage of the fact that even in an environment full of acoustic vibrations, a person's facial skin only vibrates during speech.

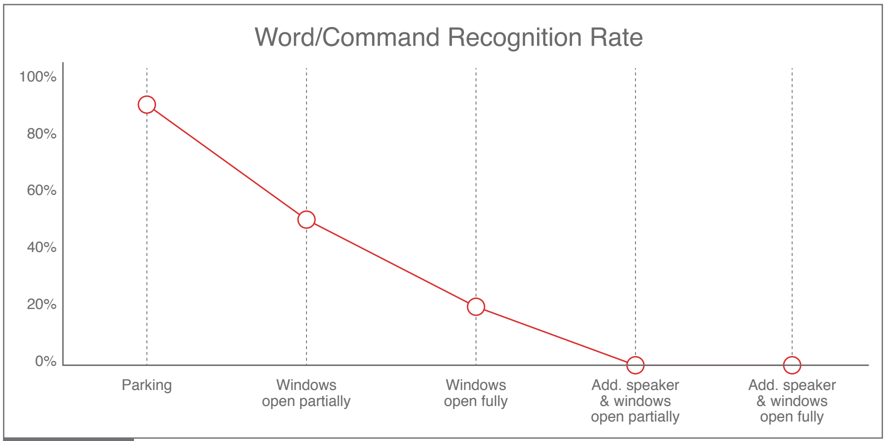

FIGURE 1 Today's voice control solutions cannot deliver what machines need to understand humans.

The optical sensor and acoustic microphone operate alongside each other. The acoustic microphone extracts signals from the air across the full 4-6 KHz range of normal speech, albeit with a high level of non-speaker-related ambient noise. Meanwhile, an eye-safe optical sensor is pointed at a fixed location on the user's face such as the mouth, lip, cheek, throat, or behind the ear, and picks up only the signals from the facial skin that are transmitted during speech at lower, 1-2 KHz frequencies (Figure 2). It is impervious to noise in this range. Nanometer-resolution interferometry techniques are then used to measure differences in the distance traveled by light as it reflects from these areas. The data is converted into intensity variations, and algorithms filter out any vibrations not associated with the user's speech. The intensity variations are then converted to signals, which are converted back to sound.

In essence, the optical HMC sensor creates a virtual "cube" around the speaker. Because vibrations are associated only with the user's speech, there is an extremely high level of directional pickup and, in turn, near-perfect isolation from extraneous noise and other background voices. No other sounds are detected or sent to the speech recognition engine.

Implementation options

The first implementation option is to connect an HMC optical sensor to the noise reduction section of a voice control solution. This improves noise reduction performance with

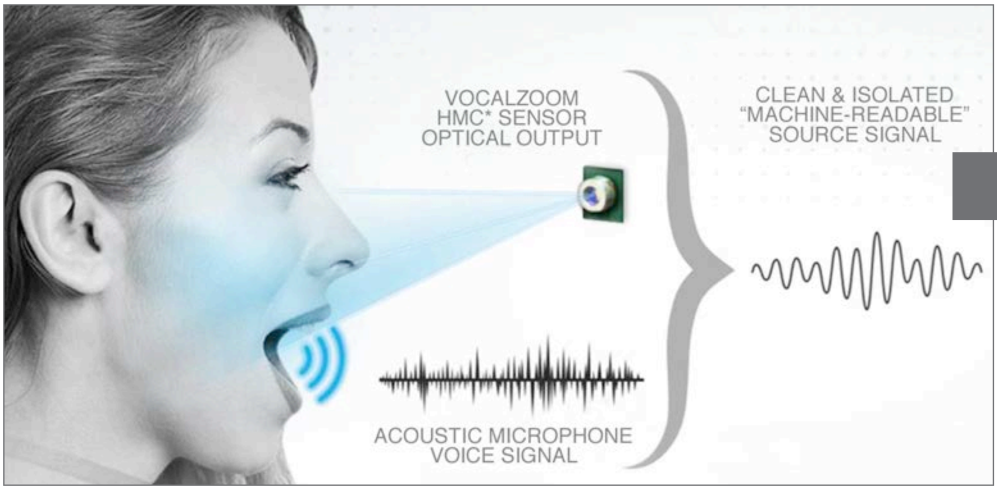

FIGURE 2

Using a multi-sensor approach, a noisy audio signal can sampled by an acoustic mic while an HMC sensor measures facial skin vibrations created by the speaker at nanometer resolution.

an associated improvement in speech-recognition performance, creating a platform for significantly improving current products without requiring changes to existing speech recognition architectures.

Alternatively, HMC optical sensors can be connected directly to a speech recognition engine, eliminating the need for noise reduction modification. The speech recognition engine simultaneously processes the acoustic and optical signals and performs all necessary noise compensation using both sets of input.

Using the latter approach, a speech recognition engine leverages the best characteristics of the acoustic microphone and optical sensor. This change to the speech recognition model has interesting implications for not only improving voice control performance but also exploring new use cases in environments that were previously considered prohibitively noisy. Today's sensor technology is small enough (sub-3 mm form factor) with sufficient power efficiency for use in very small devices, for both head-mounted (virtual and augmented reality glasses, headsets, helmets) and remote (voice-controlled automotive infotainment and access control) applications.

The connected car use case is particularly compelling. Speech recognition has been nearly impossible with windows rolled down and background passenger conversations. With optical HMC sensor technology installed in the infotainment center or rear-view mirror and pointed at the driver, however, all commands are clean and isolated from background noise. The sensor in this application operates at ranges up to 1 m across a field of view that enables typical driver movements.

Speech recognition in head-mounted devices has also been difficult, especially in noisy environments. Adding an optical HMC sensor to the headset isolates the speaker from ambient noise and removes the requirement for acoustic mics to be positioned close to the user's mouth. Designers can "remove the boom" and create new, more convenient designs and a better user experience in applications including emergency response communications solutions, motorcycle helmets, aviation headphones, and gaming and virtual/augmented reality gear. Optical HMC sensors used in these applications support an up-to-50 mm range when pointed at a fixed location on a user's face.

"HMC OPTICAL SENSORS CAN BE CONNECTED DIRECTLY TO A SPEECH RECOGNITION ENGINE, ELIMINATING THE NEED FOR NOISE REDUCTION MODIFICATION ..."

New ways to measure performance

In addition to changing how ASR engines operate and creating new voice-control use cases, optical HMC laser technology is also poised to change how the speech-recognition industry measures performance. In the past, performance was typically calculated using a Mean Opinion Score (MOS) that measures intelligibility and whether the experience is a good or bad one from a human user's perspective. The MOS has been used for decades in the telephony industry to measure quality based on a user's assessment.

In the HMC world, however, it may be more important to know how many times a command must be given before execution. Early developers of HMC solutions are now looking at such metrics as how much time it takes for a single task to be performed – i.e., using speech recognition to identify the barcode on a box on the factory floor so that an automated transport system can move it from one location to another.

Future developments

As HMC optical sensors move to the higher end of their available frequency range it will be possible to achieve unlimited vocabulary speech recognition, independent of the acoustic microphone.

Another opportunity is to point two or more optical sensors at different locations on the speaker's face, such as behind the ear and on the jaw in a head-mounted application. Add this to acoustic microphones with noise-cancellation and beam-forming capabilities, and speech recognition engines can benefit from an unprecedented level of speaker isolation for HMC applications plus ultra-high-quality audio for human-to-human communication.

It will also be possible to use a single optical HMC sensor for multiple high-value functions. For instance, optical sensors can perform proximity sensing, touch sensing (which would eliminate the need for buttons on wearable devices), and always-on voice-trigger functions. They also can be used to turn voice into another authentication factor for an expanding range of personalized online and mobile financial, healthcare, smart home automation, and other secure cloud-based services. Implementing a sensor in this way would enable system developers to replace from $10 to $20 in sensors with a single solution that leverages sensor-based interferometry for numerous applications.

The industry is moving into a new generation of capabilities with the Internet of Things and an increasingly connected world. Voice is the optimal UX, but isn't feasible without dramatic improvements. HMC optical sensor technology provides an important new solution while also creating opportunities for many new voice control applications moving forward. ECD

Rammy Bahalul, *Vice President, Sales and Business Development, VocalZoom.*

FPGAs AND AUDIO PROCESSORS ENABLE UNIQUE INDUSTRIAL APPLICATIONS

By Ted Marena

Since their inception, field-programmable gate arrays (FPGAs) were often targeted at the largest market segment – the communications industry. Although the major FPGA developers are still focused on communication applications, they are increasingly more focused on the storage and server markets.

But what about the broad industrial market?

Often, the broad industrial market's requirements are not as performance-oriented or complex as what is required in storage, server, or communications applications. For all the hype around the industrial Internet of Things (IoT) market, it is not clear how engineers can leverage available technology to make its potential a reality. An example of one technology that is seeing growing adoption in the industrial market is audio processing. By pairing the capabilities of an audio processor with the flexibility of an FPGA, many innovative applications can be supported.

An audio processor is just as the name implies – a processor that is optimized to process sound. It often leverages an ARM- based or RTOS-friendly processor architecture, has hard blocks like digital to analog converters (DACs), multiple digital microphone inputs, hardware accelerators optimized for the audible spectrum, and an I2S or SPI interface. An audio processor is normally bundled with software or firmware designed to perform certain echo cancellation or noise reduction functions.

An FPGA uses a gate-based architecture that is ideal for processing signals in a parallel fashion. It also has internal memory, hard multipliers and accumulators, and ample I/O flexibility. Some FPGAs are considered SoCs because they have quad-core, A-class ARM processors, but this level of horsepower is

not necessary when the FPGA is paired with an audio processor. The ideal pairing for an audio processor is with a generic, flexible FPGA, or an FPGA that incorporates a microcontroller such as an ARM Cortex-M3. Combining an audio processor with an FPGA of this type, with or without a Cortex-M3, creates an ideal division of labor for many tasks in unique industrial communications and control applications.

Audio listening for smart cameras

One interesting function that an audio processor can perform is audio detection when using two microphones. For example, with the appropriate firmware in an audio processor, the device can determine degree information associated with sound location. If voices or sounds are heard in front of the device (for example, in the due north position), this would be processed as "90 degrees." If sound was sourced to the right at the northeast location, the audio processor would output "45 degrees." In addition, the firmware can create a beam in front of the two microphones to attenuate noise source outside of the beam. The beam can be steered in the direction of the interested sound source by leveraging a FPGA.

This type of surveillance camera includes the following major components, also displayed in Figure 1:

1. An image sensor to capture the picture
2. An image signal processor (ISP) to handle the video data
3. An audio processor to clean up the sound path and determine location of the audio source
4. An FPGA to connect to the audio processor and implement a motor control algorithm to turn the camera towards the sound

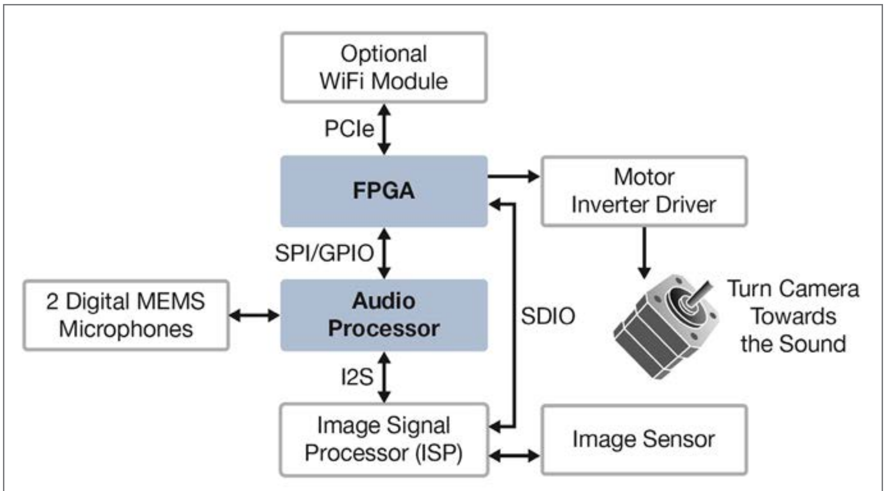

FIGURE 1 An image sensor, image signal processor (ISP), audio processor, and FPGA form the basis for intelligent surveillance cameras capable of steering towards sound sources.

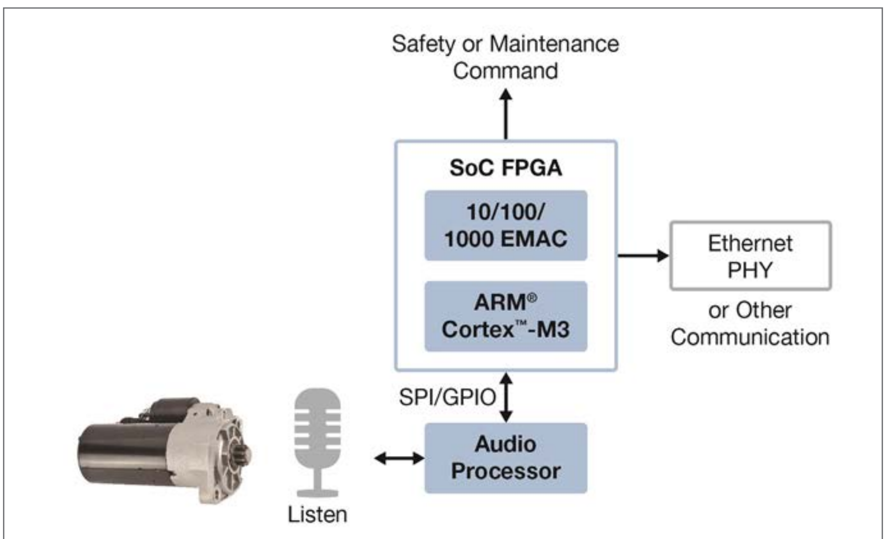

FIGURE 2 Audio processors and FPGAs can also be combined in predictive maintenance applications.

This type of smart camera could focus more closely on the image where the sound is originating. The FPGA could also be used to bridge to a high-speed Wi-Fi module based on PCIe to stream the image, or it could trigger some warning or alarm function. This approach could also be used to steer a camera and a listening microphone to optimize the performance of a videoconferencing unit. In this application, an audio processor would be used to detect where the sound originates by using the beamformer to listen to who is speaking. Instead of pointing the camera at the sound, the beamformer, under FPGA control, would also be directed to the person speaking.

> **"ANOTHER UNIQUE INDUSTRIAL IoT APPLICATION ENABLED BY AN FPGA PAIRED WITH AN AUDIO PROCESSOR IS MAINTENANCE, DIAGNOSTICS, AND FAILURE PREVENTION ... BY KNOWING THE SOUND PROFILE OF A WEAKENING MOTOR OR DRILL AND MONITORING FOR THIS AUDIO SIGNATURE WITH AN AUDIO PROCESSOR AND FPGA, PRODUCT FAILURES AND DOWN TIME CAN BE PREVENTED."**

Industrial IoT sound detection applications

Another unique industrial IoT application enabled by an FPGA paired with an audio processor is maintenance, diagnostics, and failure prevention (Figure 2 on previous page). Imagine that you had sensitive enough hearing to know when a motor or other moving component was getting weak and about to fail. Relevant examples in the industrial IoT include an elevator motor or an earth-drilling bit. By knowing the sound profile of a weakening motor or drill and monitoring for this audio signature with an audio processor and FPGA, product failures and down time can be prevented.

To implement this type of solution, sound profile firmware is inserted in the audio processor to enable monitoring of the audio signature associated with impending failure. For a motor application, this signature might be the whining sound of a bearing starting to break down or a drill whose audio profile changes to a higher level pitch as it becomes duller and has to work harder. With the sound profile residing in the audio processor, the solution then listens and continually matches the sound against the stored failure profile. Meanwhile, the FPGA talks to the audio processor and communicates to a network or some other peripheral to relay the status. If the FPGA includes an ARM Cortex-M3, it can run a lightweight TCP/IP stack and send information over Ethernet or a wireless standard. Of course, other unique communications capabilities could be leveraged such as CAN bus, USB, or a proprietary protocol.

When the audio processor detects the sound profile characteristic of a weakening condition, it signals to the FPGA and then immediately communicates this information over the network. By catching the failing condition early, the FPGA can also be programmed to trigger a response in the form of a system override. With the example of an elevator, the FPGA could wait until it secures confirmation that the elevator has reached the ground floor and everyone is out, and then communicate to the central control system that the elevator is now out of service. Additional examples where this could be used include automotive electric motors, fluid pipelines, and other industrial IoT applications. If the sound profile is well understood, this approach could even be used to trigger a call to maintenance so downtime would be significantly minimized.

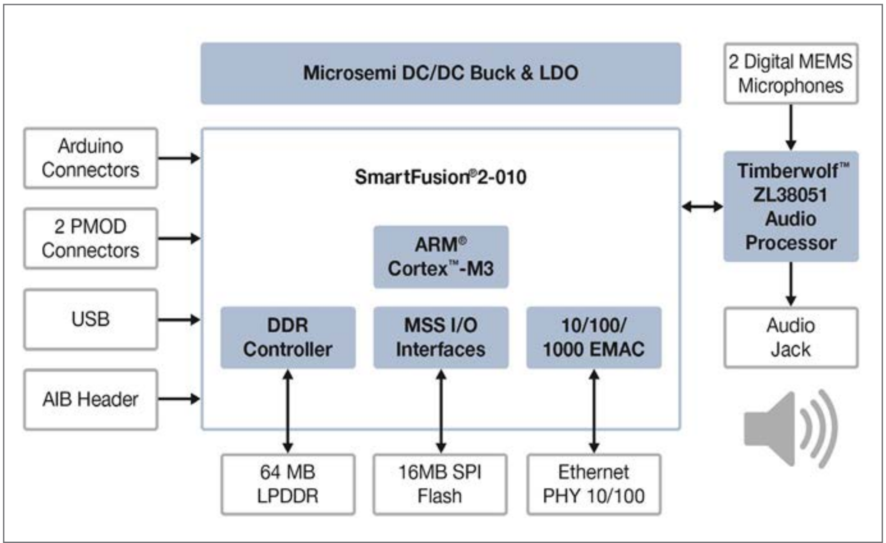

FIGURE 3 The Microsemi SmartFusion2+ (SF2+) from Arrow Electronics is an evaluation kit that combines an SoC FPGA and audio processor with reference designs for unique audio-based control applications.

Local audio storage and playback

Solutions pairing an audio processor and FPGA are also useful in applications where it is most effective to implement audio storage and retrieval locally. This includes home automation applications or to secure encrypted audio.

In these examples, the audio processor receives the voice or sound and passes it to the FPGA via an I2S bus. The FPGA then formats the data for storage in an SPI flash or other non-volatile memory. This design would also allow playback from SPI flash through the FPGA to the audio processor. Other options for this type of design include encrypting and decrypting audio for security applications. Alternatively, the FPGA could facilitate communications so the audio would be available remotely.

Recently, Arrow Electronics created a hardware kit to demonstrate the flexibility that an audio processor and an ARM Cortex-M3 FPGA can provide (Figure 3). The Microsemi SmartFusion2+ (SF2+) evaluation kit features a Microsemi Timberwolf audio processor and the SmartFusion2 SoC FPGA. The kit has on-board flash and DDR memory, as well as USB and Ethernet interfaces. In addition, a number of peripheral options can be added by leveraging the Arduino shield connector set and the PMOD interfaces.

Arrow has created a complete HDL and C code reference design for the kit that allows up to four different audio recordings to be stored, and supports playback control. The firmware for both the Timberwolf audio processor and the HDL and C code for the FPGA provide a starting point for exploring solutions that combine an audio processor and a FPGA.

Programmable control through sound

There are a number of unique and compelling applications that can be ideally implemented using the combination of an audio processor and an FPGA. The audio processor performs the task of detecting or listening for an event while the FPGA is used to provide custom responses. Of course, the additional FPGA logic also allows for custom function or other logic requirements such as bridging, hardware acceleration, or protocol communication, all of which can be explored with the availability of a hardware solution, reference design, and audio software. ECD

__Ted Marena__ is Director of SoC/FPGA Products at Microsemi.

MIXING C AND JAVA PROGRAMMING IN EMBEDDED, IOT DESIGNS

By Vincent Perrier

Although the Java language is the number one programming language in the world [1, 2], one may think that its adoption is lagging in "traditional" embedded systems because of its "fat and slow" reputation.

Actually, Java technologies changed the game in one particular type of embedded system, which is the cell phone. Cell phones have always had their specific hardware platforms and operating systems (for example, the Nokia Symbian and BlackBerry OS), but the advent of smartphones contributed to the emergence of app ecosystems such as Google's Android. Android apps are programmed in the Java language, so Java, in fact, has already won a significant percentage of embedded systems development – yes, smartphones have become "big" devices with powerful processors and plenty of memory/storage, but they're still embedded devices.

Today, the Java language is winning more and more designs in traditional non-mobile embedded systems and, in conjunction with real-time operating systems (RTOSs) and traditional C programming, is poised to become the solution of choice for IoT developers. To understand why, let's explore in more detail.

Taking embedded development to the next level

The Internet of Things is the next level for embedded systems, as IoT can be seen as "embedded" on a much larger scale:

- **Programmability:** Billions of IoT devices cannot be programmed with the limited number (in the hundred thousand range) of embedded/C/RTOS experts in the world. Industry needs to leverage larger communities (millions) of programmers from mobile/PC/server to meet the massive demands of the Internet of Things.
- **Connectivity:** IoT involves multiple wired and wireless physical layers and IP-based transport layer protocols such as UDP, TCP/IP, HTTP, TLS, REST, as well as new protocols and frameworks like CoAP, MQTT, and LWM2M.
- **Complexity:** IoT devices embed larger software content with more features and the capability to add new features dynamically (in the field) to address evolving technical or market needs.
- **User experience:** Consumers expect to interact with IoT devices as they do with their smartphones and tablets.

- **Security:** IoT devices need security at all levels – code execution, communications, identification/authentication, data storage, etc.

Java platforms provide a good solution to these challenges, as:

- Java is the number one language in the world, and all software engineering students learn it at the university level
- Java platforms offer generic implementation and APIs for IP-based networking, IoT protocols, and most non-IP protocols.
- The Java language and object-oriented programming (OOP) is well known for minimizing complexity, improving productivity, and reducing bugs.
- Java platforms enable dynamic downloading of code.
- Java platforms provide built-in security.

The key to success of Java platform implementations in embedded systems relies on tight integration with the underlying world of C, and leveraging it to the fullest extent. Java programming is not meant to replace C programming, as the C language and RTOSs are very good at providing a base runtime on top of embedded microprocessors (MPUs) and microcontrollers (MCU), and solving challenges associated with hardware-dependent software. However, Java programming is better at dealing with (developing, debugging, and maintaining) larger software packages and complexity, and at addressing hardware-independent application code.

Just like Android's virtual machine sits on top of Linux, an embedded Java platform can sit on top of an embedded RTOS and C runtime. The embedded Java platform has to be open and integrated as an independent piece of software by the C developer responsible for software bring-up on the embedded hardware, but this combined approach allows embedded projects to benefit from the best of both worlds: C for hardware interfacing and performance and Java for portability and scalability. Projects can also solve device programmability and software productivity issues as a few low-level C developers can enable dozens of higher level Java developers to build Java platforms on top of their C runtime.

THE KEY TO SUCCESS OF JAVA PLATFORM IMPLEMENTATIONS IN EMBEDDED SYSTEMS RELIES ON TIGHT INTEGRATION WITH THE UNDERLYING WORLD OF C, AND LEVERAGING IT TO THE FULLEST EXTENT.

Four key ingredients for Java integration

Java source code is compiled into a specific format called bytecode stored in class (.class) files. Class files are usually packaged into Java archive (.jar) files, which are in fact zip files that first require inflating before their bytecode can be executed. Standard Java platforms on PCs dynamically interpret bytecode with a Java virtual machine and compile it to machine code on the fly for performance improvement using a just-in-time (JIT) compiler. Unfortunately, this process cannot be transposed to MCU-based systems because it requires a lot of memory and fast processors (for storage, the inflating program, and running the JIT compiler) that are beyond the capabilities of that class of device.

But four key ingredients exist that make Java platforms suitable for integration with an embedded C-based environment with minimal memory footprint overhead (tens of kilobytes) and equivalent performance (yes, Java code can run as fast as C code). Let's review them:

1. A single, standards-based binary code format

The Executable and Linkable Format (ELF) [3] has become the de facto industry standard binary format for compiled code on MCUs. It is supported by the open source GNU GCC toochain and by other commercial toolchains. ARM, the industry-leading MCU architecture, defines its application binary interface (ABI) and relocations based on ELF.

ELF should be used as the unique and final binary code format for all programming languages used in an embedded software project.

2. Minimal onboard runtime linking

The bytecode format should not be considered as an embedded binary format, but rather as an intermediate format between the source code and the binary (machine-specific) code that is compiled and linked off-board (cross-compilation process). Off-board bytecode compilation, or ahead-of-time (AOT) compilation and linking, allows one to leverage desktop compiler optimization techniques and take advantage of the underlying instruction set and its characteristics to produce efficient code.

The Java code has to be programmed and linked into flash memory at the same time as the C code. With such an implementation, no special Java linking program is required in on-board flash: the embedded virtual machine library is just a small runtime engine that can cost only a few tens of kilobytes. All code can be directly executed in place to ensure short boot time.

3. A single, standards-based native linker

The main idea behind successful integration of a Java platform on MCU-based systems is to simply see the Java language as another programming language in addition to the C language, without having to change production toolchains used today by

C developers. This involves converting Java bytecode into ELF that can be mixed with ELF coming from compiled C code using off-the-shelf linker tools:

- The bytecode is compiled into a regular object file by a dedicated off-board compiler. Java functions are compiled to regular ELF sections, targeted by an ELF symbol with a naming convention that ensures standard ELF linkers can resolve Java symbols.
- The virtual machine is just a new ELF library added to the global project.
- The virtual machine APIs are described using regular C header files. APIs must be as generic as possible to enable porting the virtual machine to any underlying C runtime and associated RTOS, drivers, board support package (BSP), and C libraries. In extreme cases, only a timer is required when the virtual machine integrates its own internal scheduler, thus no RTOS is required.
- The whole (mixed) object files are statically linked with an off-the-shelf ELF linker. C developers still can use their favorite toolchain and integrated development environment (IDE).

Figure 1 shows the full mixed C and Java code compilation and linking steps.

4. Optimized Java-to-C code programming bridges

The embedded Java programming environment must offer access to some embedded specifics that can be done with C code:

- Immutable data (read-only data) for managing persistent (const) data stored in flash
- Bridges between Java and C programs linked as standard function calls so that any routines can be turned into C/assembly code if needed with zero-link runtime cost (linking Java code to C code is done by the off-the-shelf ELF linker)
- Fixed-size buffer sharing without any copy

The embedded Java runtime environment has to be implemented in an optimized way on top of the C runtime in order to:

- Provide an autonomous scheduler with built-in threads to ensure predictable scheduling adapted to embedded constraints ("green thread" integration to the RTOS: all Java threads run inside a single RTOS thread)
- Support object-oriented specifics (e.g., late binding in order to manage polymorphism)
- Manage memory (e.g., garbage collection adapted to embedded constraints, optimized array copy based on the C memcpy)

This enables easy reuse of legacy C code and integration of that code into the global Java application code.

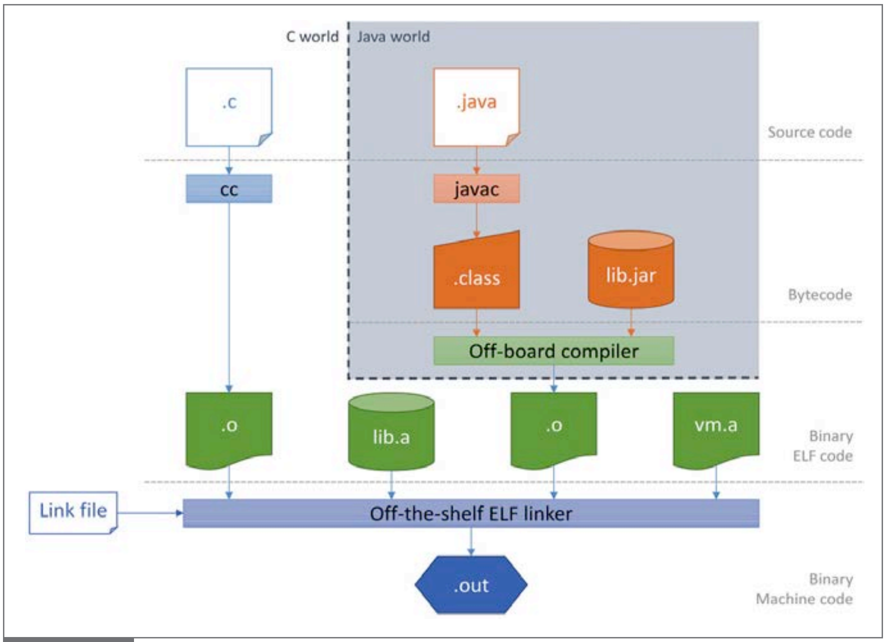

FIGURE 1 Depicted here are the code compilation and ELF linking steps in a hybrid Java and C environment.

DETAILS OF ELF

There are two main notions in ELF: symbols (names) and sections (memory zones with data or code), where basically symbols point to sections.

A symbol is an entity composed of a name and a value. A symbol may be absolute (also called a link-time constant) or relative to a section – its value will be resolved only when the linker has assigned a definitive position to the target section. A symbol can be local to the relocatable file or global to the link process. All global symbol names must be unique in the system (the name is the key to connect an unresolved symbol reference to a symbol definition).

Sections can be of two sorts:
- Allocation sections, representing a part of the program (image or runtime)
- Control sections, containing metadata (relocation sections, symbol tables, debug sections, etc.)

An allocation section can hold some image binary bytes (assembler instruction and raw data: the PROGBITS section) or can declare a runtime memory (statics, main stack, heap, etc.: the NOBITS section). A section has a conventional name representing the kind of data it holds: .text sections for binary instructions, .rodata sections for constant data, .bss sections for zero-initialized read/write data, .data sections for pre-initialized read/write data.

A relocation section is often associated with an allocation section, which contains instructions to resolve dependencies to external sections, such as a call to another function.

It is common practice in software engineering to link object files with the linker provided by the same toolchain used for compiling the object files. This allows avoiding issues when trying to link objects with different binary formats. This rule remains true for embedded C and Java programming on MCU-based systems.

The four ingredients detailed previously ensure that using the Java language for programming MCU-based embedded systems does not result in large-footprint overhead. Furthermore, developers can benefit from the compactness of the Java bytecode. Developers can use widespread Java APIs (e.g., for networking, file systems) that make software truly portable. They don't need to port their source code to heterogeneous C APIs, stacks, and compilers, or work around an unequal level of support for standards like POSIX across MCU/RTOS/compilers. Off-the-shelf binary components can be created and reused across multiple MCU architectures and associated C runtimes without porting or even re-compiling source code. Binary components can be configured at link time (using link-time constants), avoiding source-level configurations with C #define statements and interdependent source files. ECD

Vincent Perrier is Chief Product Officer at MicroEJ.

References:

1. The PYPL PopularitY of Programming Language Index: http://pypl.github.io/PYPL.html
2. The TIOBE Programming Community Index: http://www.tiobe.com/tiobe_index?page=index
3. The industry-accepted ELF specification document is chapter "Object File Format" of the "Sun Solaris Linkers and Libraries Guide" (https://docs.oracle.com/cd/E26502_01/pdf/E26507.pdf)

DETAILS ON ELF LINKING PROCESS

The linking process can be divided in three main steps:

1. Symbols and sections resolution – Starting from root symbols and root sections, the linker embeds all sections targeted by symbols and all symbols referred by sections. This process is transitive while new symbols and/or sections are found. At the end of this step, the linker may stop and output errors (unresolved symbols, duplicate symbols, etc.).
2. Memory positioning – Sections are laid out in memory ranges according to memory layout instructions described in the *linker file* (sometimes called a scatter file). Then relocation instructions are performed (i.e. symbol values are resolved and section content is modified). At the end of this step, linker may stop and output errors (if it could not resolve constraints, such as not enough memory, etc.).
3. Output ELF executable file generation – The executable file generation is associated to a memory map file which is a text file that lists what content has been linked, where it has been positioned, sizes, etc.

FEEL "FREE" TO USE EMBEDDED LINUX

By Chris Simmonds

Linux has been a mainstream embedded operating system (OS) for many years. And yet, the licensing and development model of open-source software is still little understood. Let me explain.

I will start with that that word "free," which has two meanings – free as in "you're free to use it," and free as in "nothing to pay." Linux is free in both senses, so it's a win-win. The software engineers are happy because they get access to a large body of robust, mature source code, and the product managers are happy because there are no license fees to pay. But, of the two meanings, the first is the more important. Let's be clear: Do not choose Linux just because it'll save license fees. Choose it because it's the best solution for your project.

Next, let's consider the term "Linux." Strictly speaking, Linux is the kernel of an OS, the source code of which can be obtained from www.kernel.org. More generally though, Linux is common shorthand for all the open-source components needed to create a working system based on a Linux kernel. That includes the toolchain, the bootloader, the system libraries, the command shell, the system services (also known as daemons), the end-user applications, and so on.

So, where does all this free software come from? Open source is an ecosystem, with various players fulfilling different roles. Each has their part to play, each has different motivations, and each has different revenue streams.

First is the open-source community, composed of a loose alliance of developers from many different backgrounds but with a common motivation to write software and share it with others. Why do they do it? Partly because software engineers like writing software and partly because open-source projects generate a sense of camaraderie in which you get recognition for the contributions you make. Some work purely for the buzz of doing so, but usually the core team on a big project is funded by not-for-profit organizations (such as the Linux Foundation) or by companies with a commercial interest in the technology (like Google, Red Hat, IBM, or Oracle).

You would think that putting thousands of highly individualistic programmers together would result in chaos. But that doesn't happen because open-software projects, the successful ones, at least, are highly organized. Each project is managed by one or more maintainers who control who has commit access to the source-code repository and which changes will make it into the final product.

Looking after the Inner Penguin
www.2net.co.uk

TWITTER
@2net_software

GITHUB
github.com/csimmonds

The most obvious example is the Linux kernel, lead by Linus Torvalds, and assisted by a large band of sub-system maintainers for each critical part of the kernel. Most other project maintainers keep a low profile and aren't known outside the project's circle of developers. Nevertheless, they control some of the key components on which we depend, including the GNU Compiler Collection (GCC), the GNU Project Debugger (GDB), OpenSSL, and Apache. Contributors to open-source projects are a self-selected group of highly motivated, highly talented software engineers.

The next important group of players are the vendors of the system-on-chips (SoCs) at the heart of most embedded systems. In each case, the vendor must demonstrate that Linux runs on the platform. Consequentially, they all employ large teams of Linux kernel engineers, toolchain specialists, and graphics library developers to ensure that Linux and Android work well on their platform.

For large or specialized products you may be designing a board around an SoC. For small to medium product volumes it's cost effective to use a single-board computer (SBC) or system-on-module (SoM) rather than taking an SoC and designing from the ground up.

This brings me to the third group, original equipment manufacturers (OEMs). Once again, they often have to show that they have support for Linux/Android to sell products. They take the kernel for the SoC vendor, add support for the features they've added to the board, and supply it with a Linux OS. They typically use Yocto Project to generate the full operating system, but Debian Linux is also popular. They're often a weak link in the chain, since many of them are small companies without the resources to give full support to the OS they distribute.

A fourth group are the companies that provide commercial support for embedded Linux. They can provide off-the shelf solutions for a range of hardware from SoM and SBC vendors, and can create custom Linux builds. Examples include Mentor Graphics, Wind River, Timesys, Sysgo, and MontaVista. They may bundle in proprietary components for a license fee.

Fifth, you must have noticed that there are cheap, widely available boards, such as the Raspberry Pi, which are well supported by the community. Several are available in "industrial-grade" versions that have higher spec components and can endure wider temperature ranges than the consumer versions. In addition to the Raspberry Pi, boards hail from Beagleboard.org (based on TI Sitara and OMAP3 processors), www.Minnowboard.org (Intel Atom E38xx), and www.wandboard.org (using the NXP i.MX6). All of these, except the Raspberry Pi, are open-source hardware, meaning that the schematics and board layout files are freely available (in both senses of the word), letting you customize the board to your needs. The information for the Raspberry Pi is also available, but the SoC is not, unfortunately.

Open source has always been about giving you choice, and now you have many. Which you choose depends on the in-house skill level and the time you want to devote to it. There may not be any free lunches, but there is the freedom to choose the best solution for you. **ECD**

__Chris Simmonds__ is a freelance consultant and trainer who has been using Linux in embedded systems for over 15 years. He is the author of the book Mastering Embedded Linux Programming, and is a frequent presenter at open-source and embedded conferences. You can see some of his work on the Inner Penguin blog.

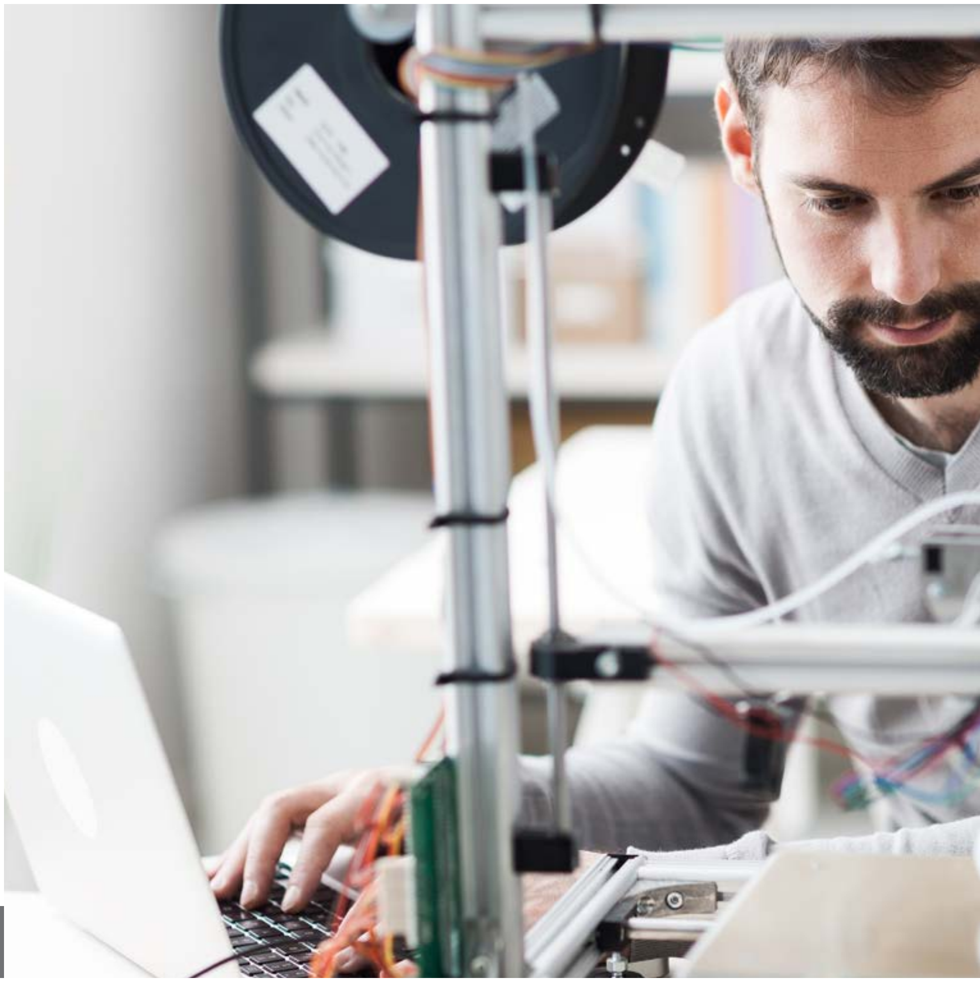

3D PRINTING EXPLAINED, AND HOW IT WILL ROCK THE WORLD

By John Hornick

3D printing is not just one process. "3D printing" and "additive manufacturing" are umbrella terms for many different technologies and processes. Each type of 3D printer builds parts or products layer upon layer, usually from the bottom up, sometimes from the top down.

3D printers have been around for about 30 years. Until recently, they were used mostly in industry for rapid prototyping. Beginning in 2009, early 3D printing patents for Material Extrusion started to expire, which led to many startups offering Material Extrusion machines (often called Fused Deposition Modeling (FDM) or Fused Filament Fabrication (FFF) machines) entering the market. This created a consumer side to the 3D printing industry. Around the same time, industrial 3D printers got good enough to start making end-use production parts, so today the industrial machines are used for both prototyping and production. Industrial machines can make end-use production parts for aerospace components, jigs and fixtures for automotive manufacturing, and customized healthcare products, such as cranial implants, surgical models, and teeth aligners.

The most common types of 3D printing technology are described below:

Binder Jetting

Also called digital part materialization (DPM), is an inkjet method somewhat like a 2D inkjet printer. Binder Jetting employs one or more jets to dispense chemical binders layer-by-layer into a bed of powdered polymer, stainless steel, bronze, tungsten (or soda lime glass powder), or sand for making molds. These machines make excellent molds and can also make finished, full-color parts after additional heat treating. 3D Systems (United States) and ExOne (United States) make industrial-grade Binder Jet printers. Voxeljet (Germany) makes Binder Jet printers with a very large build platform. A Voxeljet machine 3D printed replicas of James Bond's Aston Martin DB5 for the movie Skyfall.

Directed Energy Deposition

"Laser Engineered Netshaping" or laser cladding

Directed Energy Deposition (DED) is also known as laser cladding. Perhaps the best examples of DED systems are the Laser Engineered Netshaping (LENS) machines made by Optomec of Albuquerque, NM. LENS machines employ deposition heads, which are similar to inkjet heads, to supply metal powder to the focus of a laser beam, which melts the powder into the desired shape. Using multiple metal inputs, these machines can 3D print metal alloys on the fly. LENS machines make structural finished parts. Trumpf (Germany) has a similar DED process, and both DMG Mori Seiki (Germany/Japan) and Yamazaki Mazak (Japan) make a hybrid DED machine and multi-axis computer numerical control (CNC) mill. EFESTO makes a laser cladding machine that builds up layers of metal on existing parts. NASA and Penn State University are using the Efesto machine for a process they call Radiant Deposition, which builds up layers of metal powder radially on a rotating rod. The metal powders can be changed while a part is being built, also creating metal alloys on the fly.

Electron beam freeform fabrication

Electron beam freeform fabrication (EBFF) focuses an electron beam on metal alloy feedstock in wire form, which is fed into the beam in a vacuum, creating a molten metal pool that solidifies immediately. NASA plans to use EBFF machines to build parts in zero gravity. Sciaky (United States) calls its version of this process electron beam additive manufacturing (EBAM), and makes a giant version (9 x 4 x 5 ft.) that welds wire feedstock with an electron beam.

Material Extrusion

Fused deposition modeling (or fused filament fabrication)

Fused deposition modeling (FDM) machines extrude a thermoplastic filament, usually acrylonitrile butadiene styrene (ABS) or polylactic acid (PLA), through a tiny heated nozzle onto a build platform, building the part from the platform up. A second nozzle may extrude material to create supports that are removed after the part is built. Most consumer-level 3D printers are FDM machines, which flooded the market after early FDM patents started to expire in 2009. Currently, hundreds of companies, most of them start-ups, manufacture consumer-grade FDM machines worldwide. In the prosumer and industrial arenas, FDM machines are used mostly for prototyping, but may be used for finished plastic parts. Stratasys, an industry leader, pioneered FDM machines. Stratasys's subsidiary, MakerBot, makes consumer- and prosumer-level FDM machines, as do 3D Systems and many small companies.

Material Jetting

Aerosol Jet

Aerosol Jet machines, also called direct-write machines, use a mist generator to atomize a wide range of metal or non-metal build materials to print circuitry or parts on a variety of substrates, including onto existing parts. In this process, print-head nozzles deposit inks, such as silver nanoparticles, with extreme precision onto various substrates to make micro- and macroscale structures, such as electronic circuitry. The aerosol stream of build-material particles is refined on the fly and aerodynamically focused as it is deposited. After being deposited on the substrate, the materials may be thermally or chemically treated. Optomec (United States) is the leader in this technology. Its machines can 3D print circuitry on any substrate. Camtek's printed circuit board printer (Israel) and Neotech's light beam sintering (Germany), as well as Nano Dimension (Israel), nScrypt (United States), and XJet (Israel), seem to be using similar technology. The Lawrence Livermore National Laboratory is also working in this area.

PolyJet

Like Binder Jet machines, PolyJet machines are also multi-jet inkjet-like machines. The difference is that while Binder Jet machines jet binders onto powdered build material layer by layer, PolyJet machines jet actual build material layer by layer. Most use UV light to cure the layers of photopolymers. PolyJet machines can print multiple materials simultaneously (including support materials that are removed from the final product) and are suitable for making finished parts. PolyJet machines are made by 3D Systems, Stratasys, and Solidscape.

Powder Bed Fusion

Laser melting

Laser melting (LM) machines (also known as direct metal printing, direct metal laser sintering, metal laser melting, selective laser melting, selective laser sintering, and laserCUSING) are powder bed machines that use a laser to melt layers of plastic, ceramic, or metal powders. Because the part is fused from the surrounding bed of powder, sometimes no support structures are needed; the surrounding powder provides support, then simply falls away when the part is removed from the bed. These machines can make finished parts with complex internal and external geometries. EOS (Germany) machines use this process to make tooling and medical implants, and GE uses it to make aircraft parts, such as fuel injectors for the leading edge aircraft propulsion (LEAP) engine. Using Powder Bed Fusion, GE 3D printed, as a single piece, a cobalt-chrome fuel nozzle that formerly had been assembled by welding together twenty different parts.

3D Systems/Phenix, which calls its process direct metal printing; Concept Laser (Germany), which calls its process laserCUSING; EOS, which calls its process direct metal laser sintering; Renishaw (United Kingdom); and SLM Solutions (Germany), which calls its process selective laser melting, all make laser melting machines. Matsurra (Japan) makes a hybrid machine that combines laser melting and a CNC mill. Laser maker Fonon (United States) also makes a laser sintering machine for metal powders.

Electron beam melting

In electron beam melting (EBM) machines, an electron beam builds up parts from a powder bed in a vacuum. Similar to laser melting machines, EBM machines make finished structural parts. Sweden's Arcam is the leader here.

Sheet Lamination

Laminated object manufacturing

Laminated object manufacturing (LOM) machines laminate sheets of paper, plastic, or other materials, which are then cut into the desired shape with a laser or knife. LOM machines are well suited to making models, which feel somewhat like paper mache. MCOR (Ireland) makes LOM machines that print full-color models using standard photocopying paper. A particularly freaky example is a full-color model of an MCOR employee's disembodied head.

"ONE OF THE STRENGTHS OF 3D PRINTING IS CUSTOMIZATION. SO RATHER THAN MAKING MILLIONS OF PARTS THAT ARE ALL THE SAME, 3D PRINTING'S STRENGTH IS MAKING A MILLION PARTS THAT ARE ALL DIFFERENT. TODAY, INDUSTRIAL 3D PRINTERS ARE USED MOSTLY FOR COMPLEX OR HIGHLY CUSTOMIZED PARTS AND SMALL PRODUCTION RUNS."

Ultrasonic lamination

Another Sheet Lamination company, Fabrisonic (United States), uses a 3D printing process called ultrasonic additive manufacturing (UAM), in which sound waves fuse layers of metal foil.

Vat Photopolymerization

Digital light processing

In digital light processing (DLP), mirrors project the image of each layer of an object onto the surface of a vat of photopolymer. The light source cures the image, building up the product layer by layer. Germany's EnvisionTec and several smaller companies make DLP machines.

Stereolithography

Stereolithography (SLA) is the granddaddy of them all, the original form of 3D printing commercialized in the 1980s by Chuck Hall, who went on to form 3D Systems, an industry leader. SLA machines use a UV light source to cure a vat of liquid photopolymer resin, layer by layer. SLA-printed parts have a smooth, close-to-finished surface and are well suited for making jewelry molds. SLA can also be used for prototyping and making simple finished parts. Most SLA machines are prosumer or industrial, but consumer-level SLA machines include the Formlabs, Autodesk EMBER, and B9 SLA printers. Industrial-grade SLA machines are made by 3D Systems. The Lawrence Livermore National Laboratory is developing a high-speed variation called microstereolithography to create ultrastiff but lightweight parts.

Bringing jobs home

One of the strengths of 3D printing is customization. So rather than making millions of parts that are all the same, 3D printing's strength is making a million parts that are all different. Today, industrial 3D printers are used mostly for complex or highly customized parts and small production runs. As the machines get faster, they will make larger production runs, but the strength of the technology is mass customization, not mass production.

Because 3D printers can make entire parts or products with fewer machines, fewer steps, and therefore fewer people, they can eliminate the benefits of making things where labor is cheap. The implications are obvious: more manufacturing in America, but not many jobs running the machines. Ten manufacturing jobs lost in low-wage countries may create only one job in a 3D printing economy, but let's be careful to compare apples to apples. If it takes ten people to operate the traditional machines needed to make a single part, it may take only one person to operate the 3D printer that makes that part in America. To the optimist, that is one more manufacturing job than we had without 3D printing. To the pessimist, we still need nine more jobs. But the pessimist is missing an important point: if the part is made in America by a local worker operating the 3D printer, most of the supply, support, and distribution chain will be here too.

Regional and distributed manufacturing
Because chasing cheap labor is unnecessary in a 3D printed world, this technology can break the grip of centralized manufacturing. But don't assume that huge factories will simply replace their traditional machines with 3D printers. As 3D printers become more and more capable of making almost any finished product, centralized mass production may no longer be needed and, as a business model, may become as antiquated as the dinosaur. 3D printing will pull manufacturing away from the manufacturing hubs and redistribute it, product by product, among thousands or tens of thousands of smaller factories across the globe. Many parts and products will be made regionally, close to where they will be used.

End of the line
The days of thousands of unskilled American factory workers performing highly repetitive, mindless tasks along an assembly line are gone for good. The factory of the future will be inhabited mostly by 3D printers, robots, and other advanced machines, all driven by software. Some people will be needed on the factory floor to make sure everything is humming along, but the jobs they will do may not exist today.

As technology advances, there will be little place on the factory floor for unskilled workers. In fact, even today there are fewer and fewer jobs for workers without skills or a college education. Between October 2008, when the world economic crisis began, and mid-2014, the US unemployment rate hovered in the 6–10 percent range. During that same time period, the unemployment rate for college-educated workers was only about 3–5 percent.

In a 3D printed world, the demand for skilled workers will increase, but we don't know yet exactly what their jobs will be like. People will be needed at every step of the now-localized supply and distribution chain, even though their jobs will be radically different than they are today.

Think about the horse
So if 3D printing factories will not employ many people and most of the jobs will be for skilled workers, how will 3D printing spark a new industrial revolution, a manufacturing renaissance, and bring jobs home? Think about the horse.

When the horse was the main form of transportation, there were many horse-related jobs: saddle makers, blacksmiths, wagon makers, stable owners, feed suppliers, etc. When the automobile came along, most of those jobs were lost. But think of how many new jobs were created by the invention of the automobile. 3D printing has the same potential.

New businesses, new jobs
3D printing will spawn businesses, products, services, and jobs that are as unimaginable today as the auto industry was at the dawn of the twentieth century. Of course my crystal ball is not perfect, but some types of 3D printing-related jobs are suggested by its strengths.

Regional manufacturing means most players will be independent fabricators. A growing number of 3D printing fabricators can be found throughout the world. 3D printing fabricators are the regional and distributed manufacturers of the 3D printing age. They are the employers of the factory workers of the 3D printing–fueled manufacturing renaissance. Individually, they may not employ a large number of people, but together they will be a major source of factory jobs.

3D printing, regulation, and the American national pastime
Because 3D printing will have profound effects on stakeholders – companies, consumers, governments, and economies – it is bound to rock the law. Intellectual property (IP) law is mentioned most often, but the legal effects of 3D printing will be much broader. 3D printing will certainly affect IP law, but challenges to product safety and product liability law will probably have more relevance to most people in a 3D printed world.

Government regulators will also be challenged. Healthcare regulators will be faced with approving countless 3D printed medical devices, drugs, and human organs. Aviation regulators will face the same issues with 3D printed aircraft parts. Consumer products regulators will grapple with the safety of 3D printed products. 3D printing may also challenge governments' abilities to collect income and sales taxes, and to control the export of technology that may be used for nefarious purposes. 3D printing new kinds of crime will challenge law enforcement, investigation, intelligence, military, national security, and criminal justice systems. It will lead to calls for new laws to address the dark side.

Americans have two national pastimes. One is baseball and the other is suing each other. As companies and people are negatively affected by 3D printing, they will complain, then sue. Some will try to get Congress to enact new laws to protect their interests. Others will look for creative, positive approaches and trust in the free market system. It would be nice to think that the stakeholders will work out their problems in proactive, creative, and amicable ways. But it is more likely that 3D printing will be as much of a full-employment program for lawyers as the Internet has been. ECD

***John Hornick** has been a counselor and litigator for the Finnegan IP law firm for over 30 years. He is the author of the new book, 3D Printing Will Rock the World, and advises clients about how 3D printing may affect their businesses.*

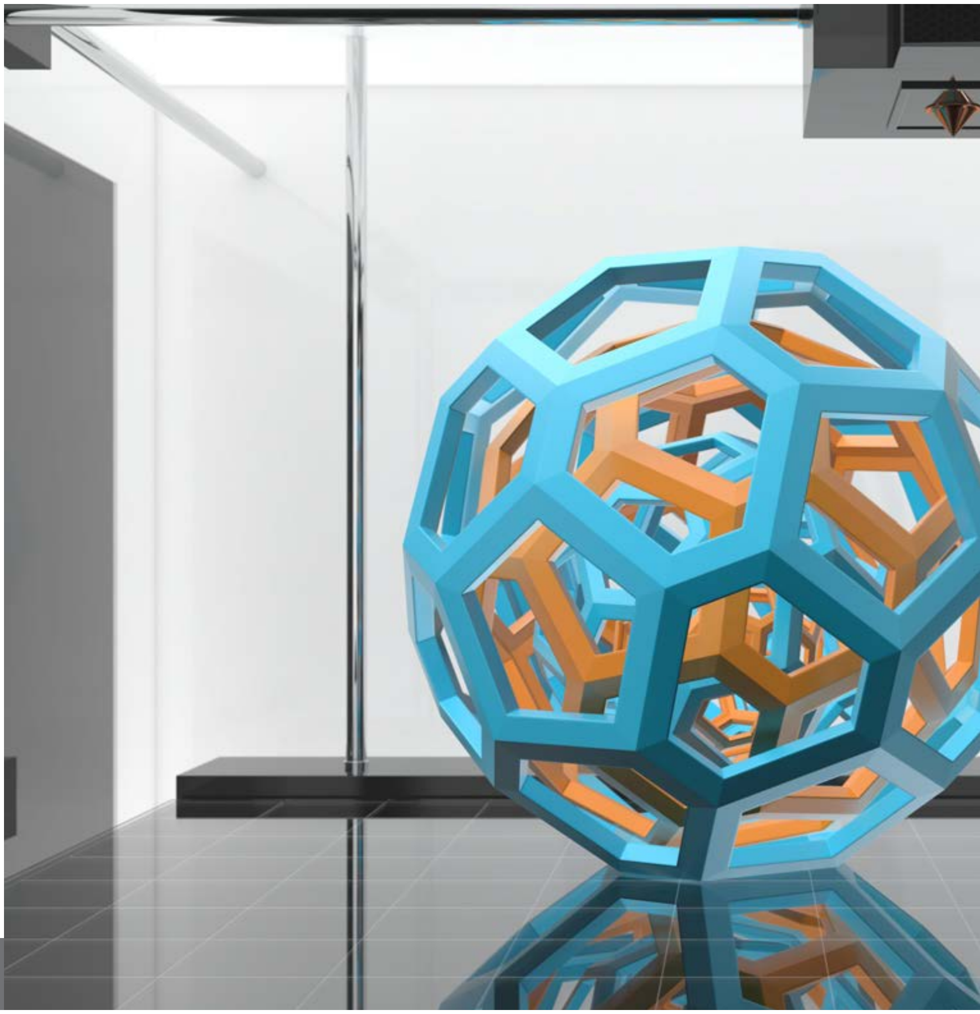

OPEN SOURCE OPENS MANY LICENSING ISSUES FOR 3D PRINTING

By Maya M. Eckstein and Eric J. Hanson

The use of additive manufacturing – commonly referred to as 3D printing – by manufacturing companies, retailers, and others is rising exponentially. PwC's April 2016 report, "3D Printing Comes of Age in US Industrial Manufacturing[1]" confirms that 71 percent of manufacturers already have adopted 3D printing and that 52 percent expect to use it for high-volume production in the next 3-5 years.

As with many technologies, the rise of 3D printing has seen a rise in the development of open source software and hardware. The use of open source brings with it many benefits, but also many pitfalls for the unwary. Users must be sure to understand the rights and responsibilities they undertake when using open source technology.

Open source software

Open source licensing often is found in a variety of software applications, printer controller software, and content files.

Software applications

3D printing requires design and modeling software applications to create object files for printing. A variety of computer-aided design (CAD), drawing, viewing, scanning, and similar content creation applications are available under open source licenses to create 3D printer-compatible files, such as the stereolithography (STL), additive manufacturing (AMF), and object (OBJ) file formats. Such software typically is released under versions of the GNU General Public License (GPL) for free use, distribution, and modification, so long as any modifications are automatically similarly licensed for "free" public use under the GPL. Other open source licenses or custom-created variations of the GPL may include commercial use restrictions, such as limiting software uses and modification to educational or personal uses. Moreover, more complex applications may include a number of software components with different licenses applicable to each component. As a result, it is necessary to review the particular open source license applicable to each software component that might be used to understand the requirements and restrictions that apply.

Hunton & Williams LLP
www.hunton.com

TWITTER
@hunton_williams

LINKEDIN
www.linkedin.com/company/hunton-and-williams-llp

YOU TUBE
www.youtube.com/channel/UCxn87ADwItC38Hrw3XuTeyA

Controller software

3D printer controller software resides on a 3D printer's microcontroller. The controller software interprets a 3D printer file and controls the hardware to print the object described by that file. While many printing manufacturers use proprietary software to control the hardware, printing materials, software, and file types used in connection with their respective printers, open source software, like that used with the Arduino and Raspberry Pi microcontrollers, have spawned new 3D printing hardware and software technologies. In fact, Microsoft recently made Windows 10 available with open source libraries to connect Arduino and Raspberry Pi microcontroller boards to Windows 10 devices. Arduino's open source licenses include the GPL for software, GNU Lesser General Public License (LGPL) for C/C++ microcontroller libraries, and Creative Commons Attribution Share-Alike license for design files. Raspberry Pi-related software is also subject to the GPL and Creative Commons licenses, as well as other third-party licenses. Because open source license terms can vary across different licenses applicable to particular open source controller software, libraries, and related software applications, it is important to consult the particular license requirements and restrictions for the applicable software component, as well as components that interact together.

Content files

Millions of design files for 3D printing are freely shared online. Such files, typically in STL format, may include, for example, original models and scanned objects. 3D content file sharing websites, such as Thingiverse, YouMagine, Pinshape, and MyMiniFactory provide platforms where users can upload and download open source software files ranging from tools and equipment to games and sculptures. On Thingiverse, users uploading files select what license is applicable to the file available for download, often the GPL or one of the Creative Commons licenses. Similarly, much of the content on YouMagine is subject to Creative Commons licenses.

A number of files on these platforms also include "custom" open source license terms that are variations of licenses like Creative Commons. Because a user sharing their design file chooses which open source license applies to that file, shared files on the platform may have different license terms. For example, some files and printed objects are restricted to only non-commercial personal uses while others permit unrestricted uses subject to providing attribution to the creator, others permit use but not sales, others dedicate the copyright interests to the public domain, and others provide different combinations of obligations and restrictions. Thus, a user downloading a shared 3D printer file cannot assume a standard "open source" license applies to all of the content files available online, but will need to review the particular license linked to the file to determine if intended uses will comply with applicable open source license terms.

Open source hardware

While open source software has become largely ubiquitous, open source hardware is less familiar. As part of the open source movement, though, open source hardware has begun to attract public attention for its potential to reduce the cost of manufacturing various goods and reshape the manufacturing and supply chain.

Open source hardware refers to physical artifacts whose design is made open to the public so that anyone can use, build, modify, distribute, and sell it free of charge. According to the Open Source Hardware Association, open source hardware "uses readily-available components and materials, standard processes, open infrastructure, unrestricted content, and open-source design tools to maximize the ability of individuals to make and use hardware.[2]" Open source hardware is intended to give people freedom to control their technology while sharing knowledge and encouraging commerce through the open exchange of designs. Open source hardware is shared by providing the bill of materials, schematics, assembly instructions, and procedures needed to make a replica of the original.

In the 3D printing space, open source hardware primarily has been used in consumer-grade printers rather than industrial-grade printers. Open source-based, consumer-grade 3D printers are now quite common, allowing for low-cost desktop 3D printers that can be used both at home and at research and educational institutions. The availability of open source-based 3D printers has helped drive the recent significant growth of 3D printing.

The largest class of open-source 3D printers are self-replicating rapid prototypers, often referred to as "RepRaps." These printers can print parts for themselves and, thus, replicate. They can even print most of the parts needed to make another RepRap. A variety of RepRap printers are available, each subject to any one of various open source licenses, including versions of the GPL and Creative Commons licenses.

Other desktop printers using open source hardware also are available. These include the Lulzbot (a product line of Aleph Objects, Inc.), published under the GPL and Creative Commons licenses; the Ember, whose manufacturer, Autodesk, touts as "the first open source, production quality 3D printer[3]" and is published under the GPL; and the Ultimaker (though it is unclear what open source license governs it).

Interestingly, one of the largest early players in the open source 3D printer market is no longer in that market. Until just a few years ago, all of MakerBot Industries' 3D printers used open source hardware. With the introduction of its Replicator 2 3D printer, though, the company opted not to share the way the physical machine is designed (or its graphical user interface) because widespread copying made it difficult for the company to succeed financially.

Indeed, the continuing use of open source hardware in 3D printers is in flux, in part because of the dilemma encountered by MakerBot and in part because burgeoning competition in the space understandably is resulting in companies seeking to protect their intellectual property. For example, 3D Systems sued FormLabs for patent infringement in 2012, accusing FormLabs of infringing eight patents. The litigation settled in 2014 with FormLabs reportedly paying 3D Systems an 8 percent royalty on all sales for a specific amount of time. While the FormLabs printer was not an open source printer, the lawsuit nevertheless unnerved the open source community. Indeed, the Electronic Frontier Foundation issued a statement in October 2012 asking for help challenging patent applications relating to 3D printers.

USES OF OPEN SOURCE SOFTWARE IN 3D PRINTING PRESENT POTENTIAL CONTRACT AND INTELLECTUAL PROPERTY INFRINGEMENT ISSUES. CONTRACT ISSUES ARISE FROM OBLIGATIONS AND RESTRICTIONS THAT MAY BE IMPOSED BY THE OPEN SOURCE SOFTWARE AGREEMENTS. INTELLECTUAL PROPERTY INFRINGEMENT CONTAINS THE RISK THAT SOFTWARE APPLICATIONS AND CONTENT FILES MAY INCLUDE MATERIAL SUBJECT TO THIRD-PARTY COPYRIGHT RIGHTS THAT IS COPIED WITHOUT AUTHORIZATION INTO THE SOFTWARE OR A 3D DESIGN FILE.

Issues

Uses of open source software in 3D printing present potential contract and intellectual property infringement issues. Contract issues arise from obligations and restrictions that may be imposed by the open source software agreements. Intellectual property infringement contains the risk that software applications and content files may include material subject to third-party copyright rights that is copied without authorization into the software or a 3D design file. In addition to copyright infringement risks, third-party trademark and patent rights might also be incorporated into 3D design files and the ultimate printed object without authorization from the rights owner.

Compliance with open source licenses

Open source licenses may include restrictions on how software and content files may be used. Some license agreements do not permit commercial uses of the licensed software. Other agreements are even more restrictive in allowing only non-commercial, personal uses that prohibit both general commercial and educational uses. Failure to review the license terms against the intended uses could lead to a breach of the license and potential damages and/or losing the ability to continue to use the open source software. Because many open source communities seek to ensure that license terms are not breached by end user entities and individuals, it is recommended that the open source license agreement be carefully reviewed for compliance.

In addition to use restrictions, many open source licenses require that the licensed software, including any modifications made by the user, be made freely available to any other users for use on the same license terms as the user was provided in the original open source software. Primarily impacting entities and users who adopt open source software for commercial uses, it is not uncommon for such users to be surprised to discover that the software applications they developed from open source software must be made freely available to others. Often this surprise comes to light when the entity user is seeking capital from a financial institution or during a business acquisition and must disclose all open source software for a consideration of how the value of the company might be impacted if the license terms permit anyone to freely copy the software. It is important for open source software users to consider whether there are "free" distribution requirements and if this impacts the user's intended goals in using such software.

Copyright infringement – Software applications and content files

Copyright infringement is probably the most prevalent risk in the use of open source software and design/content files. Because it is difficult to know who and what permission might have been provided in connection with modified source code of open source software, there is virtually always a risk that such software could include, whether intentionally or not, material copied from a third party who did not license the copyright in such material.

Such risks are even greater with content files, since third-party design files or the underlying objects may have been copied and included in the content file without the copyright owner's permission. Examples of such potential infringement include using files that contain third party characters, art works, logos, or designs that are not licensed by the third party as open source content. It is highly recommended that where material likely subject to third-party proprietary rights is the subject of a content file that confirmation of the underlying licensing of such materials be made to avoid possible infringement claims.

Copyright, patent, and trademark infringement – Printed products

Like the copyright risks that can accompany the use of open source 3D content files, the actual printing of the products from the files can lead to a number of other intellectual property infringement risks. First, the object could be a copyright infringement if it includes material copied (e.g., artwork, character) without permission of the underlying copyright owner. Second, the object could be the subject of a design patent and constitute design patent infringement absent a license from the patent holder. Third, the object (or combining several printed objects into an apparatus) could be the subject of a utility patent (e.g., a medical device or machine part) and pose a patent infringement risk if the patentee never granted a patent license. Finally, the printed object could include trademarks that are not licensed for the printed object.

In most instances, such risks are apparent when well-known third-party intellectual property rights are obvious in the content file and object. However, it is important to consider whether third-party infringement risks should be investigated as to the particular uses of 3D software and printed objects available under open source licenses. ECD

Maya Eckstein, JD, *is a partner at Hunton & Williams LLP, and heads the firm's Intellectual Property Practice Group.*

Eric J. Hanson, JD, *is a partner at Hunton & Williams LLP who focuses on intellectual property, with particular emphasis on the management of patents, trademarks, copyrights, and trade secrets.*

References:

1. "3D Printing Comes of Age in US Industrial Manufacturing." PwC. Accessed June 27, 2016. http://www.pwc.com/us/en/industrial-products/3d-printing-comes-of-age.html.
2. "Definition (English)." Open Source Hardware Association. 2012. Accessed June 27, 2016. http://www.oshwa.org/definition/.
3. "Materials for Your Business." Precision Desktop 3D Printer. Accessed June 27, 2016. https://ember.autodesk.com/.

2016 TOP EMBEDDED INNOVATORS Embedded COMPUTING DESIGN

2016 TOP EMBEDDED INNOVATOR

TYSON TUTTLE

CEO, SILICON LABS

Silicon Labs began its evolution towards the Internet of Things (IoT) in 2010, at which time Tyson Tuttle, who started with the company as a design engineer, was CTO. Since taking over as CEO in 2012, Tuttle has grown that early IoT investment into half of Silicon Labs' nearly $650 million annual revenue, and has his sights set on increasing it by driving simplicity (ironically, through software) deeper into the IoT marketplace.

A 2016 Top Embedded Innovator, Tyson Tuttle, CEO of Silicon Labs, discusses system on chip (SoC) integration for the IoT, the need to abstract complexity from embedded development, and what it's like to lead a company from the ground up.

How are SoC architectures evolving as we enter IoT era?

TUTTLE: If you look at SoCs in a cell phone, you have a big digital baseband chip, you have an LTE chip, then you've got separate power management, you've got separate radios, power amplifiers, the memory, the RAM, and the flash – all of those are a separate die. In IoT, all of those are getting integrated into a single die. So you can think of the cost point for mobile phone-type of applications being anywhere from $20 to $100 worth of electronics. In the IoT, that's perhaps a factor of 10 lower because the volumes are so much bigger but the integration level is higher. Maybe the applications are simpler, but that drives single-chip implementations.

The SoCs in a cell phone are high-powered quad-core or eight-core ARM processors, and in the IoT it might be a couple of cores worth of smaller scale ARM processors. But it's not just about processing power. It's about the integration of the connectivity, of the power management, of the sensor interfaces, of the memory, etc. So that leads you to a different process technology because you're not in a digital-only process – you have flash memory integrated in with logic, in with RF, and mixed-signal interfaces as well. So it is a different type of SoC for IoT compared to computing or mobile phones.

I would also add that the power levels are substantially different. In a mobile phone you've got a 3 amp-hour battery that lasts a day (on a good day), whereas in the IoT you may have a 200 or 300 milliamp-hour battery. So you have one-tenth the battery capacity, but that needs to last 5 to 10 years. You've got orders of magnitude difference in terms of available energy, and that fundamentally means that you're going to need to be sleeping most of the time – your chip's not going to be doing full-scale processing and communicating 24/7. You've got different energy requirements and constraints, which also leads to different optimization of the SoCs. They need to be able to hibernate and use very, very little power, and even when they're powered up they need to be able to use as little energy as possible.

For a company named Silicon Labs, how is the IoT driving your investment in software and middleware? How do you see the Silicon Labs portfolio evolving over the next 5-10 years?

"I ALSO ALWAYS TRY TO LEARN FROM EXPERIENCE AND FROM PEOPLE WITH KNOWLEDGE FROM OTHER COMPANIES WHO'VE HAD OTHER EXPERIENCES."

— TYSON TUTTLE

TUTTLE: The silicon is the foundation of the house. It is the table stakes, and ultimately that is going to be the heart of IoT devices. But then the software, stacks, and middleware are becoming a much, much bigger part of the product. You have to have a communication protocol, you have to have development tools, you have to have applications that run on those, you have to know how to communicate, and all of that is defined by software, middleware, and development tools.

Over the next 5 to 10 years the software component of our products will continue to get more and more sophisticated, so you can have a common piece of hardware that then proliferates into a lot of different versions of products that are targeting all these different applications. And, when you're talking about IoT, you're talking about thousands of different applications. So, you will have a few different varieties of hardware and different levels of optimization, but it's much more granular on the software side. We view that becoming very, very critical. If you look at the last three or four years, most of the investments in R&D and acquisitions have been, essentially, software. So, in some ways, the silicon is not the hard part.

The biggest challenges, and another aspect that's driving our investments, is this whole concept of simplicity. If you think about thousands of applications and tens of thousands of customers, being able to abstract the complexity to make it easy for customers to design in, easy for us to support, and things as simple as possible, is actually a really difficult engineering problem. But it's really the key to scaling the business in these broad markets. You can only handhold so many customers, so making things easy (a lot of which is software, some of which is also modules), building that support ecosystem, and driving simplicity is a key piece of our investment strategy in IoT over that timeframe.

You worked your way up through the ranks at Silicon Labs to eventually become CEO. How do you think that path differs from CEOs that, for better or worse, are hired from the outside?

TUTTLE: I've been with the company 19 years and started as a design engineer and worked my way up through product marketing. I started a couple of the businesses and moved into the general manager ranks. I was CTO for a while. I was COO for a while. At a technical level, a business level, and with knowledge of the teams, I've helped craft the strategy. I was not a founder, but I was the tenth employee, and have been involved in pretty much everything since day one. That is different than somebody coming from the outside.

I also always try to learn from experience and from people with knowledge from other companies who've had other experiences. Through the acquisitions we've made we've hired a lot of extremely talented people over the years, and being able to pull all that together into cohesive strategy and cohesive culture is something that has been a 19 year journey. Certainly people from the outside can come in and do a great job. Necip Sayiner, our prior CEO, came in from Bell Labs and did a great job.

So I think it can work both ways – each personality is different and each company is different – but given the culture and given my path, I think I've been able to steer the strategy in a way that builds consensus and builds enthusiasm that's different than if somebody came in from the outside and dictated. It was a little bit more from the bottom up than driven from the top down.

Silicon Labs
www.silabs.com

TWITTER
@siliconlabs

LINKEDIN
www.linkedin.com/company/silicon-labs

FACEBOOK
www.facebook.com/siliconlabs

YOU TUBE
www.youtube.com/user/ViralSilabs

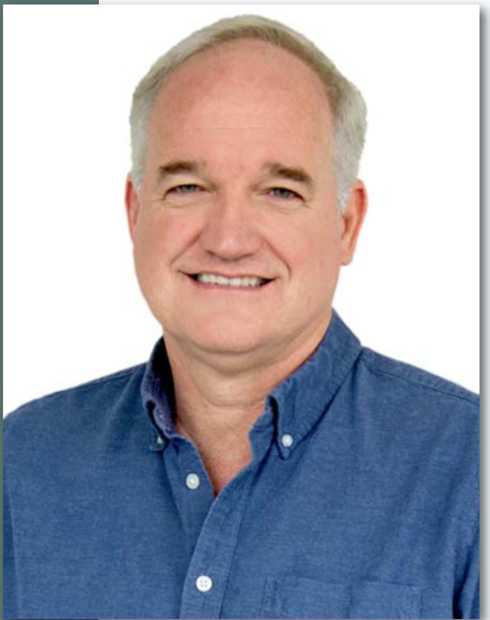

2016 TOP EMBEDDED INNOVATOR

KENT MEYER
CEO, EMCRAFT SYSTEMS

There is a disconnect between the education K-12 students are receiving and today's high-tech economies, leaving younger generations in a precarious position as technology-driven automation promises to transform the workplace (and available jobs) of the 20th and early 21st centuries.

A 2016 Top Embedded Innovator, Kent Meyer, CEO of Emcraft Systems, is rewriting the textbook with a program that puts technology instruction into the hands of industry through an Internet of Things (IoT) Educational Platform that combines kids' interest in games such as Minecraft and toys like drones with pragmatic open source technologies. Who stands to benefit? Communities, schools, industry, and the next generation of engineers.

Why is technology education falling short in schools today?

MEYER: To teach anything of substance in technology requires a considerable amount of skill, and the schools are staffed with educators who are focused on teaching the skills that tests are going to be covering – Common Core kinds of things. Getting the expertise to teachers that allows them to teach technology that is relevant, or beyond the simplest stuff, is difficult to do.

Also, when you have a class of 40 kids that you're trying to teach a science or technology topic, what you'll find is that the range of attention spans and abilities is so extreme that, if it's not self-directed, the top ones are held back and the bottom ones struggle. That's why teachers sometimes hate the science part of classes that are hands-on because it's so taxing to get everybody through a basic experiment in those numbers. It's just chaos for a teacher.

I've come to the personal conclusion that it shouldn't necessarily be the job of the K-12 schools to prepare kids for the technology jobs of the future because I think it's an impossible task. Something outside of the schools has to be provided so that kids with the interest, are able. With sports, for instance, the number of professional athletes that come back to teach kids pitching or soccer is amazing. I think the same thing has to happen for kids with technology. We do see a little of that with Brain-STEM and some of the other outfits popping up where kids can go and learn some programming classes, but those are few and far between, it's not always clear what the kids are learning, and they're very expensive programs to run. $30-$40 an hour for a kid to get some programming education is a very expensive endeavor for many parents.

All in all it's a very big problem, and one that we in the technology field are going to have to try to help the schools solve if kids are going to have the chance at something meaningful in their first 17 years as far as computing instruction is concerned.

You've developed an intriguing, industry-led education program for 6-12th graders. Can you explain how it evolved?

MEYER: The program I have running today came into existence through a long chain of events. The effort started with FIRST LEGO League (FLL) robotics for fourth graders over six years ago, because when your own kids do LEGO robotics and you're in the business of embedded microcontrollers, you get tapped to be the robotics coach.

Working with FLL there were certainly good parts of it and there were certainly frustrating parts of it where I didn't feel as though the ratio of learning to time spent was very good. So we came up with our own little curriculum where the robotics could be used to teach interested kids in a very productive way, while also trying to find entrepreneurial ways to improve the ratio of students to technology so that instead of having one robot for every ten students (which is frustrating for all) we could get as close to a one-to-one ratio as possible.

Given that and given that we had a highly motivated set of young kids who kept coming back each year, they eventually started asking for more and more things to do. So we started adding more to the curriculum. We started seeing that Minecraft could be used as a great way to teach Linux, servers, networking, and protocols, and then it just kind of ballooned. We were trying to do all of this cheaply, so we made an investment in a whole bunch of off-lease Chromebooks and started running Linux in developer mode on those, then we started teaching Arduino.

That's how we've come to where we are. The kids keep growing up and they keep getting smarter, so the things that you need to teach them are getting more complex.

How has the IoT contributed to the success of these efforts?

MEYER: I view the IoT as one of the best gifts an educator of technology and software could have because it lets you go from the smallest, tiniest devices all the way up to the cloud, and it provides a framework that allows kids to understand how things are working together. That's always been one of my complaints about any kind of class that you'd teach because a Python or JavaScript class is cool and you learn some great stuff about computer science, semantics, and syntax, but it's usually a one-off type of course and then you move on to the next thing without it fitting into a larger framework.

Really what we're defining here is the IoT Educational Platform. If we use the concepts that were successful with LEGO robotics and the Minecraft classes where we leverage what the kids are already interested in and apply it around an IoT Educational Platform, kids can work on projects and continually add things by working up and down the IoT stack. Now you have something that is a living, breathing platform that the kids are contributing to a framework that they can annotate based on what they're learning and understand how the various pieces interact. The key idea is how to make something the kids can build upon, from the very beginning with the simplest classes to the more complex, and how all those parts can work together.

What do programs like yours centered need to get further off the ground?

MEYER: A goal I have for the kids this summer is to have the Carlsbad Mayor, City Council, and the school board come to our shop and see what these kids are doing so we can present the idea of a Makerspace for kids focused on Industry 4.0 and IoT. If we can get the City of Carlsbad and the school district to come and see what these kids are doing with some new projects – the younger ones are going to try to simulate Amazon delivery through Minecraft using drones; we've got a guy making an electric skateboard; and another is going to be working on database integration and analytics in the cloud and trying to tie into IBM Bluemix – my honest to god hope is that we can convince the City of Carlsbad to set up an incubator-type facility.

Because the schools can't each individually afford to set up a high-tech learning center at each of the 11 different K-12 schools here, the idea is to instead make some focused investments. If we could set up a facility with a well-defined curriculum that has the backing of the school district, the city for the facility, and then industry as far as people donating time or resources, there are a number of high-tech companies that could really benefit from having kids learn skills at an early. The requirements analysis could be done at the beginning so that we can say to industry, 'Hey, if a student has these skills the summer after their junior year of high school, would you be interested in having them as an intern to help you?' I'm confident that we could get a huge number of local businesses excited, especially if you went to them and said, 'That's what you want them to know? Here's the roadmap for how we can teach them those skills at this facility.'

The whole strategy is to try to get the right people talking to get something like this set up, make the right investments, and then turn the whole thing around so that industry tells us what they're looking for. For the kids that means at least an internship, and maybe even jobs.

Emcraft Systems
www.emcraft.com

EMAIL
kent@emcraft.com

LINKEDIN
www.linkedin.com/company/emcraft-systems

GITHUB
www.github.com/EmcraftSystems

2016 TOP EMBEDDED INNOVATORS
Embedded COMPUTING DESIGN

2016 TOP EMBEDDED INNOVATOR

JEFF BADER
VICE PRESIDENT, EMBEDDED BUSINESS UNIT, MICRON TECHNOLOGY

While much has been made of the Big Data and cloud storage aspects of the Internet of Things (IoT), much less attention has been paid to the growing memory requirements of devices that exist at the edge. Here, as connectivity, security, and demands for localized intelligence increase, so too do the expectations for low cost, small footprint, and minimal power from embedded storage solutions.

A 2016 Top Embedded Innovator, Jeff Bader, Vice President, Embedded Business Unit at Micron Technology, has overseen the growth of an automotive and industrial memory portfolio that now ships more than 1 million parts per day. In this interview with Embedded Computing Design, Bader offers his take on the current memory bottlenecks for IoT edge devices, advances in memory technology and packaging, and a vision for what integrated storage solutions could look like in 5 to 10 years.

Why is storage at the edge of IoT critical, and what memory bottlenecks currently exist there?

BADER: Device storage in the IoT is a rapidly growing area for both memory and digital storage that is growing from the continued distribution of compute and communications capabilities to the edges of an ever-expanding network of connections. As the capability of those edge devices increases, the memory and storage needs within them are increasing as well. At the same time, the challenges of connectivity – bandwidth, latency, and persistence – are driving the need to increase both storage and data intelligence even further in those edge devices.

As the costs of compute and storage continue to decline, it enables the distribution of the data analytics toward the edge, which can provide faster response times, better decision support, or simpler overall system design. Ultimately, we want real-time decisions from our devices, not just data. By leveraging local storage and the intelligence to analyze, devices can provide richer information for upstream analytics engines. In terms of the memory bottlenecks or key memory criteria today, systems are really driving a balance of performance, reliability, and cost in addition to things like power, form factor and longevity, depending on the specific device.

In order to truly see the scale of IoT that has been forecast, the interoperability standards and related security concerns of edge devices need to improve. Recent studies have suggested that a significant percentage of all IoT devices today are lacking even very basic security features, and security at the system level can be significantly improved by secure storage features within memory devices.

What is being done to reduce the cost of storage for embedded/IoT devices? What wafer or process technologies are emerging to enable this?

BADER: Faster adoption of leading edge technology is occurring in embedded applications that traditionally used technology only after it had been adopted and in stable production by mainstream applications. We see this, for example, with the strong adoption of LPDDR4 technology in automotive applications today.

While we do see the adoption of leading edge lithography to enable cost reductions within embedded/IoT devices, this is balanced by a need within these applications for extended lifecycles and higher durability. As a result, memory solutions must be developed that support both cost reductions and improved durability, such as the family of industrial-grade SSDs Micron recently introduced. Another key trend driving system cost reductions is the increasing use of multi-chip packages (MCPs) and system-in-package (SiP) solutions. Integrating DRAM with non-volatile memory in an MCP or combined with an SoC in an SiP solution can deliver significant system cost savings in addition to performance and power benefits.

5 to 10 years from now, what is going to be the storage architecture of choice for Industrial IoT and embedded systems?

BADER: The crystal ball gets a little foggy 10 years from now, but it's probably safe to say some of the same key trends we see today will continue in embedded and IoT. For higher function IoT devices, the proliferation of mobile architectures into embedded systems will continue, taking advantage of the scale, performance, and power of those systems. This means advanced DRAM, eMMC, and universal flash storage (UFS) will combine in MCP and package on package (PoP) configurations to deliver the best performance and form factor possible. For lower function devices, we'll continue to see NOR as the primary boot solution with a growing use of low-density NAND and eMMC. In the longer term we could also see the adoption of emerging memory solutions like Micron's transistor-less 3D Xpoint technology, which could serve as a simple non-volatile RAM solution replacing (or augmenting) both DRAM and NAND in those devices.

Micron Technology, Inc.
www.micron.com

TWITTER
@MicronTech

LINKEDIN
www.linkedin.com/company/micron-technology

HONORABLE MENTIONS

2016 TOP EMBEDDED INNOVATORS

Skip Ashton
Vice President, IoT Software, Silicon Labs

Alicia Asín
CEO, Libelium

Ciaran Connel
CEO and Co-Founder, DecaWave

Cesare Garlati
Chief Security Strategist, prpl Foundation

Scott Hanson
Founder and CTO, Ambiq Micro

Jeffrey Hibbard
CEO, IntervalZero

Keith Jackson
President and CEO, ON Semiconductor

Raj Johnson
President and CEO, Mapusoft Technologies

Sherri McDaniel
President, ATEK Access Technologies

Robert Miller
Founder, Technologic Systems

Todd Ouska
CTO, wolfSSL

Hans Rempel
CEO and Co-Founder, Exosite

Andy Rhodes
Executive Director,
Commercial IoT Solutions, Dell

Evan Thomas
Assistant Professor, Portland State University,
and CEO, SWEETSense, Inc.

Serge Tissot
Technical Fellow and Principal Architect,
Technology Platforms, Kontron

Lorie Wigle
Vice President and General Manager,
IoT Security Solutions, Intel Security

2016 TOP INNOVATIVE PRODUCT NOMINEES

This year, the fourth installment of Embedded Computing Design's annual Innovation Issue includes a highly competitive spread of product submissions from every area of the embedded/Internet of Things (IoT) development space. As in years past, each submission will be judged based on a rubric that considers performance, features, and, most importantly, how disruptive the solution is in the market.

39	AMBIQ MICRO	INTEL QUARK TECHNOLOGY
	DIGILENT INC.	INTERSIL
	IMAGECRAFT CREATIONS INC.	
40	MARVELL TECHNOLOGY GROUP	MULTI-TECH SYSTEMS, INC.
	MAXIM INTEGRATED	NXP SEMICONDUCTORS
	MERCURY'S ENSEMBLE	
41	SANDISK CORPORATION	SIERRA MONITOR CORPORATION
	SENDYNE	SOMNIUM TECHNOLOGIES
	SEQUITUR LABS INC	
42	TAG-CONNECT	ZEBRA TECHNOLOGIES CORPORATION
	TECHNOLOGIC SYSTEMS	
	TELEDYNE LECROY	

Embedded Computing Design editors Rich Nass, Curt Schwaderer, and Brandon Lewis will vote to select the Top Innovative Products of 2016, three of which will be featured on the cover of the Embedded Computing Design Resource Guide in August.

Ambiq Micro — **Apollo MCU**

The Apollo MCU is an ideal solution for applications looking to reduce or eliminate the need for batteries and lower overall system power. The expanded power budget afforded by the Apollo enables the possibility of adding features and functions that may not have been previously possible. The Apollo MCU family, is the industry's first ever microcontroller to rely on Subthreshold Power Optimized Technology (SPOT) enabling devices to operate circuits at subthreshold voltages, while utilizing standard CMOS technology. The Apollo MCU offers industry-leading power consumption; consuming just 34µA/MHz when executing instructions from flash with sleep mode currents as low as 150nA.

ambiqmicro
The low power revolution starts here

www.youtube.com/watch?v=46rKUgUTr4k
www.ambiqmicro.com **www.embedded-computing.com/p373598**

Digilent Inc. — **Digilent Nexys Video High Definition Video Platform**

The Digilent Nexys Video is targeted for high definition video acquisition, processing, and display applications, based on Xilinx Artix 7 FPGA (XC7A200T-1SBG484C). The 512MB 800MHz DDR3 memory makes full HD video input and output stream seamlessly over the on-board mini Display or HDMI connectors. An FMC connector allows user to have high speed data communication between the I/O interface and FPGA. The video targeted reference design enables an immediate start to software, firmware and hardware development. An open source high definition video processing project (4K resolution) developed by our community member is available at Digilent forum

http://store.digilentinc.com/nexys-video-artix-7-fpga-trainer-board-for-multimedia-applications/
www.youtube.com/watch?v=bLfXhdxs_Yk **www.embedded-computing.com/p373599**

ImageCraft Creations Inc. — **Wifi2go**

With ImageCraft's wifi2go module, designing IoT products just got a whole lot easier! Wifi2go combines an STM32F4 Cortex-M4 MCU with the state-of-the-art TI CC3100 wifi chip with robust security. What makes wifi2go unique is ImageCraft's JumpStart C development environment: you aren't limited by cloud-based IDEs that weren't designed for professional embedded development. JumpStart C runs much faster than Eclipse, and with JumpStart API, the 32-bit Cortex-M is now as easy to program as an 8-bit MCU. Wifi2go is OEM-production-ready, and can be soldered directly onto PCBs. Best of all, JumpStart C starts at only $249. Faster, better, in-budget: get your IoT products to market now!

https://imagecraft.com/promos/osm-video-2016.html
www.imagecraft.com/wifi2go **www.embedded-computing.com/p373600**

Intel® Quark™ Technology — **Microcontroller Developer Kit D2000**

The Internet of Things (IoT) is fueling innovation in nearly every part of our lives. Simply connecting the "things" that were never connected before is leading to new data insights that translate into meaningful change. Intel® Quark™ technology is the onramp for all of this data making your "things" intelligent and connected. The newest Intel® Quark™ microcontroller D2000 is a low power, battery-operated, 32-bit microcontroller with a robust instruction set. It enables scalability and interoperability on one x86 architecture from the edge to the data center. The D2000 includes integrated device level security and access to Intel's vast ecosystem and development toolset. Intel®, enabling your full end-to-end IoT story.

www.intel.com/quark **www.embedded-computing.com/p373601**

Intersil — **ISL8273M**

Intersil's ISL8273M is the industry's first 80A fully encapsulated digital DC/DC PMBus power module. It's a complete step-down regulated power supply that provides point-of-load conversions for advanced FPGAs, DSPs, processors and memory. It operates from industry standard 5V or 12V input power rails and supports multi-phase current sharing of up to four ISL8273M power modules to create a 320A solution with output voltages as low as 0.6V. Housed in a compact 18mm x 23mm module, the ISL8273M delivers the industry's highest power density and performance for increasingly space-constrained data center servers and storage equipment, as well as wireless communications infrastructure systems.

www.youtube.com/watch?v=z98qOzFDPDk
www.intersil.com **www.embedded-computing.com/p373602**

Marvell Technology Group — AndromedaBox™

The highly integrated AndromedaBox™ IoT platform is designed specifically for the rapid prototyping and development of smart connected devices. It's built on Marvell's award-winning, high performance ARMADA® 385 application processor, Marvell's 88E6141 GE switch that supports 2.5 Gigabits per second uplink and Marvell's advanced Avastar® Wi-Fi 11AC, Bluetooth and 802.15.4 connectivity solution. The AndromedaBox™ platform is made for Brillo, Google's operating system, with native support for Weave communication protocol, and also supports Ubuntu Linux, OpenWrt and other open software stacks. The platform enables router and gateway functions, home automation and optional private cloud service through a NAS connection.

www.youtube.com/watch?v=pBj8eA1skO8
www.marvell.com **www.embedded-computing.com/p373603**

Maxim Integrated — MAX14720

Maxim Integrated's MAX14720 power management integrated circuit (PMIC) allows designers to optimize power and battery life for wearable medical/fitness and IoT applications. Key advantages include lower power, extended battery life, longer product shelf life, and flexible operation and system diagnostics. Increasing battery life and achieving low power are common challenges faced by engineers when developing wearable and IoT products. The MAX14720 is ideal for non-rechargeable battery applications where size and energy efficiency are critical. An electronic battery seal extends shelf life by effectively disconnecting the battery prior to initial power-up. Integrating the functionality of five discrete devices, the MAX14720 reduces the bill of materials and allows for much smaller form factor designs.

www.maximintegrated.com/en/products/power/battery-management/MAX14720.html
www.maximintegrated.com **www.embedded-computing.com/p373604**

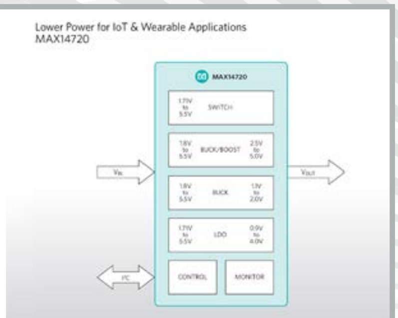

Mercury's Ensemble® — LDS6526

Mercury's Ensemble® LDS6526 is an OpenVPX processing blade that seamlessly integrates an Intel® Xeon® processor D system-on-a-chip (SoC), a versatile software-defined off-load processing and a built-in, double-bandwidth sensor I/O capability into a powerful 6U form factor blade for streaming signal processing applications. Mercury's software-defined, FPGA-based protocol offload engine technology (POET™) combined with Altera's latest Arria® 10 device deliver twice the sensor I/O bandwidth than any other OpenVPX blade with four channels of I/O that can be routed to either the processor or data plane. This innovative embedded technology is switch fabric-agnostic and runs 40Gb/s Ethernet or many other protocols.

www.mrcy.com/LDS652
www.mrcy.com **www.embedded-computing.com/p373605**

Multi-Tech Systems, Inc. — MultiConnect® xDot™

MultiConnect® xDot™ is a secure, CE/FCC certified, ARM® mbed™ programmable, low-power RF module, that provides long-range, low bit rate M2M data connectivity to sensors, industrial equipment and remote appliances. The MultiConnect xDot is LoRaWAN™ 1.0 compliant, providing bi-directional data communication up to 10 miles/15 km line-of-sight and 1-3 miles/2 km into buildings**, using sub-GHz ISM bands in North America and Europe. xDots bring intelligence, reduced complexity and a lower overall bill of material cost to the very edge of the network while supporting a variety of electronic interfaces to connect just about any "Thing" for years on battery power.

https://youtu.be/RayXSSWnVpw
www.multitech.com **www.embedded-computing.com/p373606**

NXP Semiconductors — Hexiwear

From smart health to smart home, Hexiwear transforms your great idea into a finished product by pairing open-source, optimized hardware with a stylish, compact form factor for exceptional usability and sophisticated functionality. Based on NXP technology, the Hexiwear platform is equipped with powerful MCUs, on-board integrated sensors, a color OLED display, rechargeable battery, external flash memory, and the option to add additional sensors through click boards™. Supported by a comprehensive open-source software suite that includes embedded software, cellphone apps and cloud connectivity, Hexiwear is your complete wearable design platform for the IoT era.

www.youtube.com/watch?v=HL90eduDSDw
www.nxp.com **www.embedded-computing.com/p373607**

SanDisk Corporation — AEC-Q100

Built to meet the intense reliability, quality and temperature requirements, AEC-Q100 certified SanDisk® Automotive storage solutions bring reliable, high-performance storage to variety of connected, in-vehicle applications, including 3D mapping and advanced augmented reality in navigation systems, entertainment systems, intuitive driver assist technology and more. The latest SanDisk Automotive SD card is also enhanced with smart features that offer enhanced power failure protection, memory health status monitoring, customization capabilities that enable OEMs to exclusively designate and program the card specific usage then remotely monitor the card to ensure it's operating at ideal performance or to identify when upgrades are needed.

www.sandisk.com/oem-design/automotive
www.sandisk.com **www.embedded-computing.com/p373608**

Sendyne — dtSolve™

dtSolve™ is the first model solver designed specifically for embedded/IoT applications, such as robotics, motor control, pattern recognition, etc. where model-based control (MPC) is a necessity. MPC is only as effective as the model used for that control. While models that truly represent the physics of any given system are used successfully in desktop simulations, today these same models cannot be easily utilized for embedded use. Instead they must be overly simplified, leading to a loss of accuracy and the need for expensive system overdesign to hedge against the unknown. dtSolve™ makes it possible to use actual physical models for MPC. What's more, the same dtSolve™ software runs on a desktop and on a microcontroller – the same tool can thus be used for every stage of model development and deployment.

www.sendyne.com **www.embedded-computing.com/p373609**

Sequitur Labs Inc. — CoreTEE™

CoreTEE™ is a Global Platform compliant Trusted Execution Environment for IoT and other embedded devices. CoreTEE leverages the ARM® TrustZone®. CoreTEE provides a multi-tenant, secure foundation to create trustworthy devices by providing isolation for critical processes, data and key material. CoreTEE is based on the open source OpTEE project managed by Linaro (www.linaro.org) with significant enhancements critical for commercial deployments. OEMs deploying CoreTEE gain the following benefits: Reduced BOM costs by replacing external (legacy) TPM components; Secure boot and firmware updates field-upgradeable cryptography algorithms; A scalable security platform for future applications; Accelerated time-to-market with out-of-the-box trusted applications

www.youtube.com/watch?v=UWQvCFcr1sc
www.sequiturlabs.com **www.embedded-computing.com/p373610**

SEQUITUR LABS

Sierra Monitor Corporation — Cloud-Connected IIoT Gateway

Sierra Monitor's cloud-connected IIoT Gateway is a high-performance protocol gateway that seamlessly integrates industrial devices, such as power meters and uninterrupted power supplies (UPS), into building management systems in commercial buildings and industrial facilities. The IIoT Gateway also supports a set of web applications that allow users to drill down into connected devices to see device data, trends, and events. The set of applications can be accessed locally onsite and remotely in the cloud through Sierra Monitor's FieldPoP™ Device Cloud. A user can manage all of their devices anywhere, from one web-based interface, in a secure fashion with no firewall dependencies.

www.sierramonitor.com **www.embedded-computing.com/p373611**

Somnium Technologies — DRT Cortex M-IDE

SOMNIUM® DRT Cortex M-IDE is the innovative embedded C/C++ development environment for ARM® Cortex® M microcontrollers. DRT is unique – fully compatible with industry standard software enablement and GNU tools whilst offering significantly better code generation. DRT's patented resequencing optimizations typically provide 20% smaller code, 15% higher performance and 30% energy savings with no source code changes required. DRT's IDE enhances developer productivity with automatic project import and advanced debug and trace (including live memory viewing – an essential feature for real-time embedded systems). DRT automatically provides the best possible technical results with reduced engineering time, effort and costs.

www.youtube.com/watch?v=46rKUgUTr4k
www.somniumtech.com.com **www.embedded-computing.com/p373612**

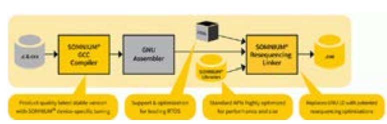

2016 TOP INNOVATIVE PRODUCT NOMINEES

Embedded COMPUTING DESIGN

Tag-Connect — TC2030-CTX (-NL) Plug-of-Nails™ Programing Cables for ARM® Cortex®

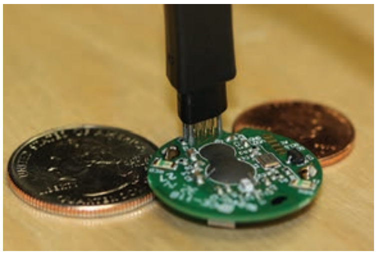

Tag-Connect's TC2030-CTX (-NL) Plug-of-Nails™ Cables plug directly into a tiny footprint of pads and holes in your PCB eliminating the traditional JTAG header. With a no-height footprint about the size of an 0805 resistor the Tag-Connector saves all of the cost and most of the space of a header making it perfect for flat or space constrained products such as wearables. Aimed at ARM Cortex applications, the TC2030-CTX cables have a regular 10-pin Cortex connector at one end and a 6-pin Plug-of-Nails™ Spring-Pin connector at the other. Seen here a TC2030-CTX-NL cable connects to EM Microelectronic's tiny BLE beacon. Tag-Connectors are seen in a large percentage of BLE (aka BS) devices ("things") such as most BLE beacons, smart locks, and wearables.

www.tag-connect.com/TC2030-CTX
www.Tag-Connect.com **www.embedded-computing.com/p373613**

Technologic Systems — TS-7680

The TS-7680 embedded computer, powered by a 450 MHz ARM CPU, is designed to provide extreme performance for applications which demand high reliability, fast bootup/startup, and consistent connectivity. The TS-7680 was built to handle power input allowing 8 to 40 VDC as well as 10 to 28 VAC. It is designed to operate fanless in the industrial temperature range of -40 °C to +85 °C.

The TS-7680 is our inaugural launch of our TS-SILO super capacitor technology providing up to 20 seconds of power for a graceful shut down in the event of power loss and to ensure file system integrity.

www.embeddedarm.com/products/TS-7680
www.embeddedarm.com **www.embedded-computing.com/p373614**

Teledyne LeCroy — ProVIDE

ProVIDE, combines two important test and debug methods – software IDE and Protocol Analysis. Initially, offered as a plugin to the Eclipse Integrated Development Environment, It extends the IDE debug capabilities to allow analysis and visualization of both internal and external busses. Capitalizing on Teledyne LeCroy's market leading protocol analyzer tools, ProVIDE gives developers a unique view into how their control and data actually travels across the external bus and synchronizes this information with the code execution. Using IDE controls and breakpoints with Teledyne LeCroy protocol analyzers, developers can isolate and correlate their code with actual signal traffic over the interface.

www.teledynelecroy.com/provide **www.embedded-computing.com/p373615**

Zebra Technologies Corporation — ARM® mbed™ OS IoT Starter Kit powered by Atmel

The wireless IoT technology preview kit provides users with a seamless and smooth experience for prototyping and getting started with IoT. The kit powered by Zatar, an ARM-mbed enabled IoT Cloud Service and Atmel's (a wholly owned subsidiary of Microchip) Xplained board enables a preview of standards-based chip-to-cloud stack for creating next generation Internet-aware devices.

The kit consists of a 32-bit ARM® Cortex®-M4 RISC processor, OLED1 Xplained pro and a WINC1500 fully certified module with Wi-Fi & network capabilities. The ability to get to market quickly is an increasingly critical design factor as the IoT opens up new possibilities.

https://youtu.be/Gpwnxu4mn0g
www.zebra.com | www.zatar.com **www.embedded-computing.com/p373616**

2016 TOP INNOVATIVE PRODUCT

Embedded COMPUTING DESIGN

THE WINNERS WILL BE ANNOUNCED IN THE UPCOMING AUGUST ISSUE

IPSec/IKEv2 boosts network security

HCC Embedded has added an IPsec/IKEv2 module to their embedded software portfolio to provide secure, authenticated communications between IPv4/IPv6 hosts and gateways. Security applications include medical device communications, point-of-sale terminals, industrial control, and machine-to-machine (M2M) applications. The IPsec/IKE implementation comes with a static analysis report based on MISRA compliance.

HCC Embedded
www.hcc-embedded.com
www.embedded-computing.com/p373596

Dedicated IoT network connectivity joins forces with smart home automation

The SIGFOX global Internet of Things network will be used by Groupe HBF for connectivity of its OTIO ALERT product line that detects domestic incidents for smoke, carbon monoxide emissions, large temperature swings, power failure, and home intrusion. The network will enable users to be immediately notified by SMS in an alert situation. The SIGFOX network is ideal for the application because of its multi-national reach and dedicated IoT network that provides energy-efficient transmission of small quantities of IoT data over large distances.

SIGFOX
www.SIGFOX.com | http://iotdesign.embedded-computing.com/news-id/?51438

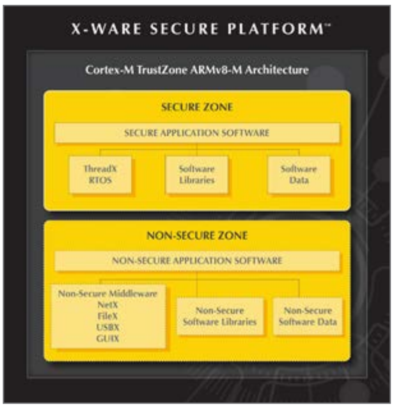

ARM-based software platform for secure information and safety-sensitive IoT

Express Logic's X-Ware Secure Platform uses ARM TrustZone technology to protect ARMv8-M and MPU-based IoT devices. The X-Ware software leverages the Cortex-M MPU to assign memory partitions that isolates non-trusted code from the RTOS and other trusted applications. This also facilitates custom application software downloads and operation of that code within these partitioned zones, as the rest of the system remains protected from possible errors in the downloaded code. In addition, the X-Ware environment provides network I/O, graphics, USB, and other middleware components needed for IoT device development.

Express Logic
www.rtos.com | www.embedded-computing.com/p373597

Innovation for Integrated Power-Telecom networks

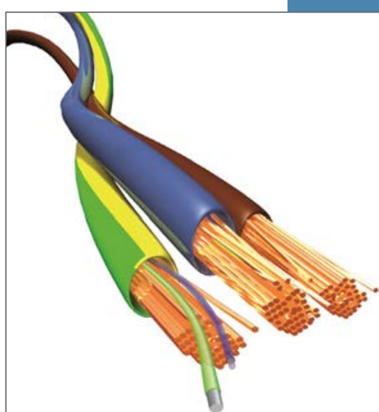

Prysmian Group announced a new cabling system able to combine power and ultra-broadband voice and data connectivity to homes and businesses. The cabling system combines cable for energy transport and embedded fiber optics for passive connectivity. The system can also include an active electronic switch to enable information gathering and analytics between power plants, power cabinets, and substations. There is even a component to provide ultra-broadband connectivity. The ability to trench and run lines once for multiple purposes represents significant cost savings while simultaneously enabling industrial IoT applications for both power grid and network services.

Prysmian Group
www.prysmiangroup.com
www.embedded-computing.com/p373508


Empowering the Connected World
Boosting IoT Designs from Edge to Cloud
▪ Intel® Core™ based module, IoT gateway, network appliance
▪ Energy-efficient embedded solution
▪ Industrial operating temperature
▪ 7+ product longevity
▪ High reliability
Enable Your Design Today
www.portwell.com
info@portwell.com
1-877-278-8899
Portwell
intel
IoT Solutions Alliance