

newelectronics

9 March 2021

DESIGNING FOR THE IOT • ADVANCED PLATFORMS • THERMAL MANAGEMENT



SMARTPHONES JUST GET SMARTER

The market for smartphones
continues to thrive

Access to 9.6 Million+ Products Online



DIGIKEY.CO.UK

9.6 MILLION+ PRODUCTS ONLINE | 1,200+ INDUSTRY-LEADING SUPPLIERS | 100% FRANCHISED DISTRIBUTOR

From Design to Production

**FREE
SHIPPING**

ON ORDERS OVER
£33 OR \$50 USD*

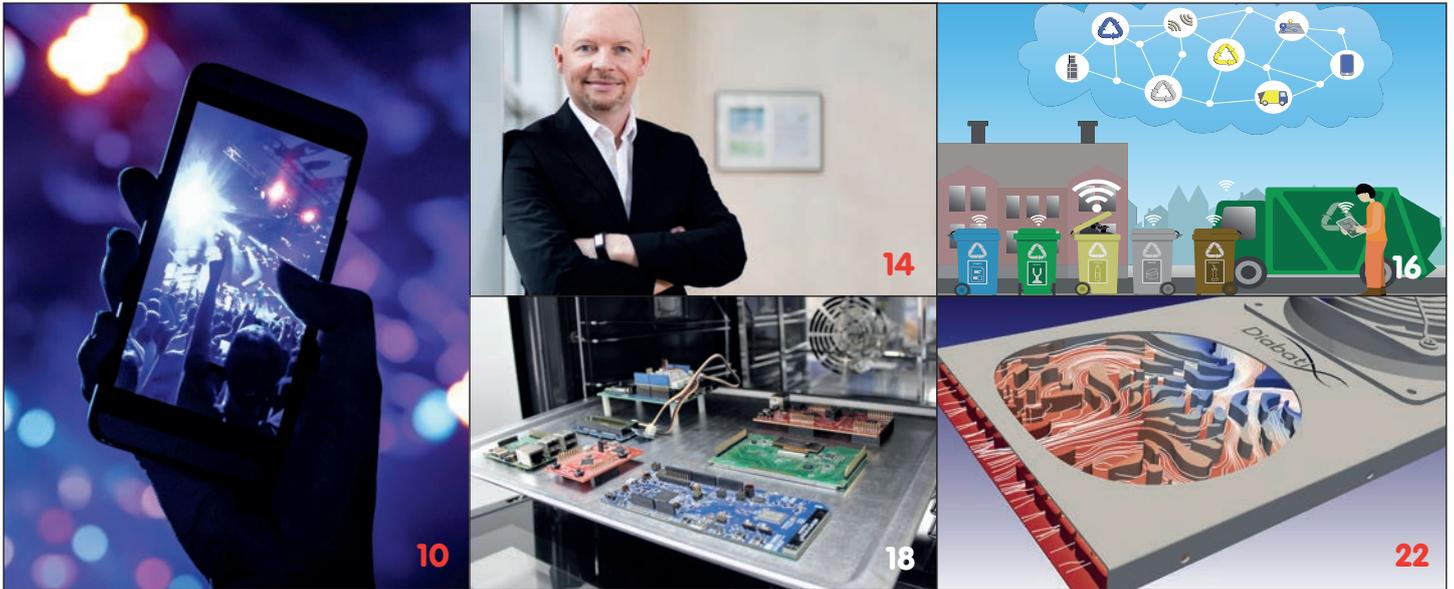


0800 587 0991
DIGIKEY.CO.UK



*A shipping charge of £12.00 will be billed on all orders of less than £33.00. A shipping charge of \$18.00 USD will be billed on all orders of less than \$50.00 USD. All orders are shipped via UPS, Federal Express, or DHL for delivery within 1-3 days (dependent on final destination). No handling fees. All prices are in British pound sterling or United States dollar. Digi-Key is a franchised distributor for all supplier partners. New products added daily. Digi-Key and Digi-Key Electronics are registered trademarks of Digi-Key Electronics in the U.S. and other countries. © 2021 Digi-Key Electronics, 701 Brooks Ave. South, Thief River Falls, MN 56701, USA

ECIA MEMBER
Supporting The Authorized Channel



COVER IMAGE: 9parusnikov/stock.adobe.com

COMMENT 5

Can research and innovation help to drive a post pandemic economic recovery?

NEWS 6

Heat-free optical switch would enable optical quantum computing chips

MediaTek looks to usher in AI-enabled interactive multimedia experiences

Skin-comparable sensor to help robots perform more challenging tasks

COVER STORY 10

Smartphones just get smarter

With the roll-out of AI accelerating and leading to the creation of new applications, the market for smartphones continues to thrive. By Neil Tyler

INTERVIEW 14

Delivering 5G's infrastructure

Brendan Farley, MD of Xilinx EMEA, discusses the phased roll-out of 5G infrastructure and the opportunities that will create. By Neil Tyler

DESIGNING FOR THE IOT 16

Finding the 'unexpected' bug

It's a real challenge to identify bugs in embedded code across the Internet of Things, but visual trace diagnostics has made that process easier. By Johan Kraft

CONNECTED PRODUCT DESIGN 18
Cooking up a connected product

Adding a connected device to your portfolio can be challenging. Jonathan Pallant looks at how to pick the 'perfect ingredients'

ADVANCED PLATFORMS 20

Addressing safety critical FPGA designs

Advanced EDA tools and methodology is helping designers to avoid synthesis bugs in safety-critical FPGA designs, as Rob van Blommestein explains

THERMAL MANAGEMENT 22
Thermal evolution

Generative design and machine learning point to novel ways to cool systems, as Chris Edwards discovers

NEXT GENERATION COMMUNICATIONS 25
A blended strategy

Why using a blend of networks will deliver a more successful Internet of Things deployment. By Nick Sacke

MISSION STATEMENT

'New Electronics keeps designers and managers abreast of the latest developments in the world's fastest moving industry'



Delivering more

The widest selection of semiconductors and electronic components in stock and ready to ship



mouser.co.uk

Reed Relays for **High Voltage** Applications

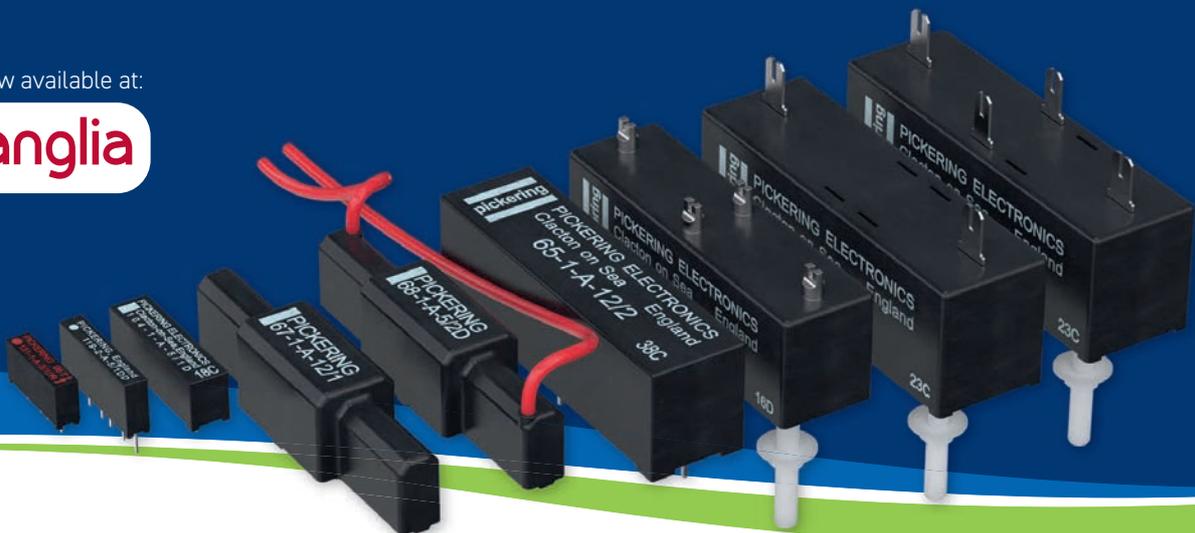
Ideal for high-end cable testers, mixed signal semiconductor testers, medical electronics, electronic vehicles, solar energy, backplane testers, in-circuit test equipment and high voltage instrumentation.

For a **free** working sample go to: pickeringrelay.com/samples

- Switching 500 V to 12 kV at up to 50 W
- Available in 1 or 2 Form A and 1 or 2 Form B, and with 5 kV Form C coming soon!

Now available at:

anglia



pickeringrelay.com
sales@pickeringrelay.com

Product
25+Years
Longevity

100%
Tested



Turning aspirations into a reality

CAN RESEARCH AND INNOVATION HELP TO DRIVE A POST PANDEMIC ECONOMIC RECOVERY?



Last week's Budget expected a swifter bounce-back from the pandemic than had previously been expected. The Chancellor, Rishi Sunak, used his Budget speech to promise "honesty" about what would be needed to rebalance the economy and claimed he was laying the foundations for a post-pandemic recovery.

Sunak talked at length about making the UK, "the best place in the world for high-growth, innovative companies," and announced a number of measures intended to encourage business investment and an R&D innovation-led recovery.

But was it anywhere near enough? For many commentators it was a thin statement, even when taking into account the headline grabbing 6 per cent increase in corporation tax.

Commenting CaSE Executive Director Professor Sarah Main said that while there had been a strong innovation theme, with support for innovative firms and the green economy, announcements on science funding were "notably scarce". She pointed to there being no provision for the UK's agreed participation in Horizon Europe, no specific measures to alleviate Covid pressure on UKRI's budget, or support for research charities – all of which will have a significant impact on the UK's science and research activity.

That's not to say that there were no announcements of note - the introduction of the Future Fund to support the scaling up of R&D-intensive businesses and the launch of the British Business Bank, to help companies access the capital they need to develop and grow, should be welcomed.

These announcements recognise the potential of high-growth, innovative technology companies to the UK and of the green industrial revolution.

Science has an important role to play in the UK's recovery and future growth.

However, as Professor Sir Jim McDonald, the President of the Royal Academy of Engineering said, "Our ambitions on net zero, infrastructure and digitalisation are threatened if we do not have the number and diversity of people with engineering and technical skills needed to deliver them."

That means we should be looking to attract high profile talent from around the world, for sure, but also at the potential of communities in UK regions outside of traditional scientific hubs.

Young people have been disproportionately affected by the pandemic and the government has really got to work at creating opportunities for them – it's no good talking about levelling up if you don't deliver on those promises.

For research and innovation to support the UK's recovery and growth it must start from a strong foundation and that only comes from a talented and educated workforce.

Neil Tyler, Editor (neil.tyler@markallengroup.com)

"Despite a strong innovation theme, with support for innovative firms and the green economy, announcements on science funding were "notably scarce".

CaSE Executive Director Professor Sarah Main

newelectronics

Editor Neil Tyler
neil.tyler@markallengroup.com

Contributing Editors Chris Edwards, Charlotte Hathway,
John Walko
editor@newelectronics.co.uk

Art Editor Chris Charles
chris.charles@markallengroup.com

Illustrator Phil Holmes

Sales Manager James Creber
james.creber@markallengroup.com

Publisher Peter Ring
peter.ring@markallengroup.com

Managing Director Jon Benson
jon.benson@markallengroup.com

Production Manager Nicki McKenna
nicki.mckenna@markallengroup.com

New Electronics editorial advisory panel

Trevor Cross, chief technology officer, Teledyne e2v

Pete Leonard, electronics design manager, Renishaw

Pete Lomas, director of engineering, Norcott Technologies

Neil Riddiford, principal electronics engineer, Cambridge Consultants

Adam Taylor, embedded systems consultant

ISSN 0047-9624 Online ISSN 2049-2316

Annual subscription (22 issues):
UK £108. Overseas; £163. Airmail; £199.
New Electronics, incorporating Electronic Equipment News
and Electronics News, is published twice monthly by
MA Business, Hawley Mill, Hawley Road, Dartford, DA2 7TJ.
T: 01322 221144 E: ne@markallengroup.com

Moving on?

If you change jobs or your company moves, please contact
circulation@markallengroup.com to continue receiving
your free copy of New Electronics



A MARK ALLEN GROUP COMPANY
www.markallengroup.com

© 2021. All rights reserved. No part of New Electronics may be reproduced or transmitted in any form, by any means, electronic or mechanical, including photocopying, recording or any information storage or retrieval system, without permission in writing from the Publisher. The views expressed do not necessarily represent those of the editor of New Electronics. Advertisements in the journal do not imply endorsement of the products or services advertised.

Please read our privacy policy, by visiting
http://privacypolicy.markallengroup.com. This will explain
how we process, use & safeguard your data
Printed by Pensord.



When you have finished with
this magazine please recycle it.

Boost for quantum computing

HEAT-FREE OPTICAL SWITCH WOULD
ENABLE OPTICAL QUANTUM COMPUTING
CHIPS. NEIL TYLER REPORTS

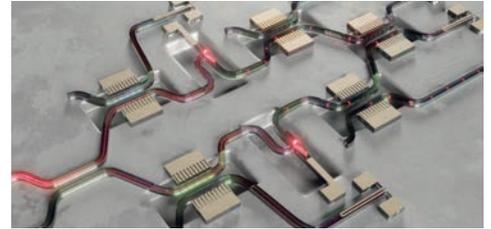
In what could be a boost for quantum computing and communication, a team of European scientists have reported a new method of controlling and manipulating single photons without generating heat.

The European Quantum Flagship project, S2QUIP, is reported to have developed an optical switch that is reconfigured with microscopic mechanical movement rather than heat, making the switch compatible with heat-sensitive single-photon detectors.

Currently, optical switches work by locally heating light guides inside a semiconductor chip. "This approach does not work for quantum optics," said Samuel Gyger, a researcher from the KTH Royal Institute of Technology in Stockholm.

"Because we want to detect every single photon, we use quantum detectors that work by measuring the heat a single photon generates when absorbed by a superconducting material. If we use traditional switches, our detectors will be flooded by heat, and won't work at all."

The new method enables control of single photons without the disadvantage of heating up a



semiconductor chip and rendering single-photon detectors useless.

By using MEMS actuation it is possible to enable optical switching and photon detection on a single semiconductor chip while maintaining the cold temperatures required by single-photon detectors.

"Our technology will help to connect all the building blocks required for integrated optical circuits for quantum technologies," said Errando Herranz, who led the work at KTH.

"Quantum technologies will enable secure message encryption and methods of computation that solve problems today's computers cannot and they will provide simulation tools that enable us to understand fundamental laws of nature, which can lead to new materials and medicines."

Herranz said that the group is looking to integrate the fabrication process in semiconductor foundries that already fabricate on-chip optics.

Micro mobility pilot project

DUCKT, a micro mobility infrastructure innovator, has been awarded a pilot project for the installation of 150 dock, lock and charge points in Paris. The project aims to demonstrate how a universal charging infrastructure could be used to accelerate micro mobility use to reduce climate impact in the city.

The charge solution will be deployed in Q2 2021 and DUCKT will be offering users a way to help organise public space, lower operational costs and provide a simple, secure universal charge station. The adaptable solution can be plugged into advertising boards, bus stations and street lighting to provide a power source.

Marc-Antoine Réol, Country Manager France at DUCKT said: "Paris was ideal to demonstrate the power of our solution. With nearly 15,000 electric scooters on the streets, policies are going to keep moving towards accelerated deployment of micro mobility. The aim of this project is to provide a last mile infrastructure, which links public transport and shared micro-mobility."

Active Silicon acquired

Active Silicon has been acquired by Solid State, a manufacturer of computing, power and communications products, and a supplier of electronic and opto-electronic components.

Operating through two main divisions, Manufacturing (Steatite) and Value Added Supplies (Solid State Supplies and Pacer), the group specialises in complex engineering often requiring design-in support and component sourcing for computing, power, communications, electronic and optoelectronic products.

The £6.3m acquisition will see the addition of Active Silicon's expertise in the design, manufacture and supply of vision components to the Solid State portfolio, and will further its operations in the areas of machine and computer vision, and provides Active Silicon with greater potential to develop their embedded computing and AI capabilities.

Active Silicon will operate as an independent entity within the industrial computer product portfolio of Steatite's manufacturing division.

MediaTek unveils 4K smart TV chip

MEDIATEK LOOKS TO USHER IN AI-ENABLED INTERACTIVE MULTIMEDIA EXPERIENCES.

NEIL TYLER REPORTS

MediaTek has announced a new 4K smart TV chip, the MT9638, which comes with an integrated high-performance AI processing unit (APU).

The MT9638 has been designed to support AI enhancement technologies such as AI super resolution, AI picture quality and AI voice assistants, plus variable refresh rate (VRR) and MEMC (motion estimation and motion compensation) so that graphics will appear smoother.

The device's high processing capabilities means that it has been optimised for visual quality, and will give global TV makers features to design competitive 4K smart TVs.

"As smart home appliances become more intelligent and diversified, smart TVs are becoming the new hub for the smart home as consumers take advantage of AI-enabled multimedia features to control all their devices right from their TV," said Alex Chen, General Manager of the TV Business Unit at MediaTek.



With built-in AI picture quality technology, the chip enables real-time content and scene recognition, automatically adjusting colour saturation, brightness, sharpness, dynamic motion compensation and smart noise reduction to improve the overall image quality.

MediaTek's AI super resolution technology, combined with its MEMC technology, makes it possible to intelligently upscale resolution through multi-frame blending to deliver clearer images at the smart TV's native resolution.

The chip also supports the latest connectivity technologies including Wi-Fi 6 and HDMI 2.1, which integrates support for VRR to match movies or console gaming frame-rates in order to avoid screen tearing.

The MT9638 also supports applications powered by AI voice recognition and AI enabled video calls, opening up new ways for consumers to interact with their smart TV.

SMART Modular launches DDR5 Module family

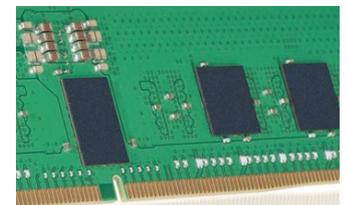
SMART Modular Technologies has launched its DDR5 module family of products, as it looks to support the transition to higher-performing, more efficient memory technologies.

DDR5 has been designed to provide increased memory bandwidth and performance to help process data across various application segments including HPC, AI, and edge computing and provides twice the performance of DDR4 with a 4.8 Gbps to 6.4 Gbps data rate and improved power efficiency over DDR4 by using an on-DIMM 12V voltage regulator (PMIC) and 1.1 I/O voltage.

Another important factor contributing to its higher performance is its dual-channel DIMM topology for higher channel efficiency that makes it suitable for growing server, cloud computing, and enterprise networking applications.

DDR5 is able to achieve increased performance by doubling the burst length to BL16 and bank-count to 32 from 16 on DDR4. For increased reliability and efficiency, DDR5 DIMMs also have two 40-bit fully independent sub-channels on the same module and have a maximum die density of 64Gb with the capability of scaling to much higher DIMM storage capacity than DDR4 DIMMs.

SMART Modular's initial DDR5 module product family includes RDIMMs, UDIMMs and SODIMMs in densities ranging from 16GB to 64GB.



Compliance Suite targets secure IoT applications

Secure Thingz, an IAR Systems Group company, has announced Compliance Suite, a set of tools and training for embedded developers.

These tools and training are specifically targeted to provide developers with a simplified path to building applications that are compliant with the European EN 303645, UK & Australian 13 Best Practices, and the evolving US Cybersecurity Improvement Act (NISTIR 8259).

Compliance is a challenge when it comes to the fast-moving Internet of Things domain so the IoT Security Foundation has developed an IoT Security Compliance Framework, enabling

organisations to build a self-certification methodology that meshes with the 13 Best Practices captured in UK and European Secure by Design guidelines.

The Compliance Suite includes a set of development tools and Preconfigured Security Contexts that enables developers to rapidly implement core aspects of the guidelines, such as moving from passwords to certificate-based identification; the implementation of update policies; and the use of

advanced device specific security enclaves to protect provisioned information.

Coupled with these tools is a set of training and support resources linking the functional requirements with the certification requirements identified in the IoT Security Foundation Compliance Questionnaire, ensuring a rapid implementation that meets international requirements.





SLAMcore offers 'out-of-the-box' robot navigation

SLAMcore, a specialist in Spatial Intelligence for autonomous robot location and mapping, has added new dense 2.5D mapping capabilities to its latest software release.

Robots require maps of their surroundings and the objects within it to navigate effectively - these maps can include point clouds, 2D flat maps, 2.5D dense maps (with heights), 3D maps and semantic maps which identify objects (people/tables/chairs/windows).

One of the main challenges of robot mapping is that it requires significant compute power, processing time and memory to create these maps.

In response, SLAMcore has developed software that uses stereo cameras and inertial sensors to build 2.5D height maps in real-time. The occupied space is represented as a series of columns of different heights which show what space is filled. Using these rich 2.5D maps, robots and autonomous devices know where objects are and can safely plot routes through real-world environments.

SLAMcore software is designed to run fast prototypes out-of-the box with Intel RealSense depth cameras (D435i or D455) and has been optimised for x86 and Nvidia Jetson processors. The software can be further customised for production systems to run on a wide range of cost-effective hardware from Raspberry Pi to GPU based systems.

Scientists develop soft tactile sensor

SENSOR WITH SKIN-COMPARABLE CHARACTERISTICS COULD HELP ROBOTS ACCOMPLISH MORE CHALLENGING TASKS.

NEIL TYLER REPORTS

A joint research team co-led by City University of Hong Kong (CityU) has developed a new soft tactile sensor with skin-comparable characteristics. According to scientists, a robotic gripper using the sensor would be able to accomplish a range of challenging tasks such as stably grasping fragile objects or threading a needle.

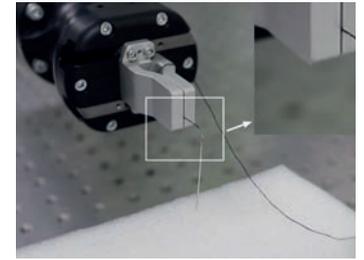
The research has provided new insight into tactile sensor design and could contribute to various applications in the robotics field, such as smart prosthetics and human-robot interaction.

A main characteristic of human skin is its ability to sense the shear force, meaning the force that makes two objects slip or slide over each other when coming into contact. By sensing the magnitude, direction and the subtle change of

shear force, our skin can act as feedback and allow us to adjust how we should hold an object stably with our hands and fingers or how tight we should grasp it.

This new sensor comes in a multi-layered structure like human skin and includes a flexible and specially magnetised film of about 0.5mm as the top layer. When an external force is exerted on it, it can detect the change of the magnetic field due to the film's deformation. More importantly, it can "decouple", or decompose, the external force automatically into two components - normal force (the force applied perpendicularly to the object) and shear force, providing the accurate measurement of these two forces respectively.

The sensor also possesses another human skin-like



characteristic - the tactile "super-resolution" that allows it to locate the stimuli's position as accurately as possible.

"We have developed an efficient tactile super-resolution algorithm using deep learning and achieved a 60-fold improvement of the localisation accuracy for contact position, which is the best among super-resolution methods reported so far," said Dr Shen Yajing, Associate Professor at CityU's Department of Biomedical Engineering (BME).

This efficient tactile super-resolution algorithm can help improve the physical resolution of a tactile sensor array with the least number of sensing units, so reducing the number of wirings and the time required for signal transmitting.

ABB launches next generation cobots

ABB has extended its collaborative robot (cobot) portfolio with the launch of the GoFa and SWIFTI cobot families, offering higher payloads and speeds to complement YuMi and Single Arm YuMi in ABB's cobot line-up.

Designed to be stronger, faster and more capable these cobots are intended to help accelerate the company's expansion into high-growth segments such as: electronics, healthcare, consumer goods, logistics and food and beverage, and meeting the growing demand for automation across multiple industries.

GoFa and SWIFTI are intuitively designed so that customers don't have to rely on in-house



programming specialists.

Sami Atiya, President of ABB's Robotics & Discrete Automation Business Area said, "These robots are easy to use and configure and they are backed by our global network of on-call, on-line service experts to ensure that businesses of all sizes and new sectors of the economy, far beyond manufacturing, can embrace

robots for the first time."

ABB's cobot portfolio expansion looks to respond to four key megatrends including individualised consumers, labour shortages, digitalisation and uncertainty that are transforming business and driving automation into new sectors of the economy.

A global survey of 1650 large and small businesses found that 84 percent said they will introduce or increase the use of robotics and automation in the next decade, while 85 per cent said the pandemic had been "game changing" for their business and industry, with it acting as a catalyst for accelerating investment in automation.

Now: VNA Mode For Real-Time Spectrum Analyzers



**Our Newcomers:
Great in Price and
Performance!**

Vector Network Analysis Mode (VNA, Standard):

- S11, S21 and Distance-to-error Measurement (DTF)
- Smith, Polar, SWR and Group Delay can be displayed

RTSA Mode (Real-time):

- Up to 40 MHz Real-time Bandwidth
- FMT, Density, PVT, Spectrogram



from **€ 7,895**
plus VAT

**3 Years
Warranty –
Extendable!**

from **€ 2,099**
plus VAT

RSA5032N / RSA5065N

- 9 kHz to 3,2 or 6,5 GHz Frequency Range

GPSA Mode (Search):

- -165 dBm (typ.) Displayed Average Noise Level (DANL)
- -108 dBc/Hz Phase Noise

RSA3015N / RSA3030N / RSA3045N

- 9 kHz to 1,5 / 3 or 4,5 GHz Frequency Range

GPSA Mode (Search):

- -161 dBm (typ.) Displayed Average Noise Level (DANL)
- -102 dBc/Hz Phase Noise

www.rigol.eu

For more information please contact your local RIGOL Partner

NEW FHD 13.3" Display with eDP interface

The TM133VDGP01 is the latest display to come out of Tianma. This 13.3" a-Si TFT has FHD of 1920 x 1080 and an eDP interface. With high brightness built-in and a colour range of 16.7M this display is ideal for factory automation, kiosks, military and transportation. It also has the advantage of a long production life cycle with an upgrade path to brighter units.

RDS is ISO 13485 Qualified to design and manufacture complete systems.

Innovative Design and Manufacture of IoT Platforms and Devices.

TIANMA

rds



www.review-displays.co.uk

+44 (0)1959 563 345
e: info@review-displays.co.uk



MANHATTAN SKYLINE

Size	Res.	P/N	Status	Kick Off	CAS	1 st Sample	Final Sample	MP
8.4"	1024 x 768	ET084030DVA	Developing	Sep. 18	Oct. 12th	End Dec.	Jan. 2021	2021
10.4"	1024 x 768	ET104030DSA	Developing	Oct. 8th	Oct. 23th	End Dec.	Jan. 2021	2021
12.1"	1024 x 768	ET121030DVA	Developing	Sep. 7th	End Sep.	End Nov.	Jan. 2021	2021

EDT releases replacement TFTs for Mitsubishi's discontinued display modules

Following Mitsubishi Electric's official press release in June 2020, where the company announced:

The company's LCD business, currently focusing on small and medium-sized TFT-LCD modules for industrial and automotive use, has been expanding its business by developing and launching high value products with advanced technologies in the specific market segments based on the market needs. However, the company has decided to end the TFT-LCD modules production and terminate the business based on the consideration that the company is no longer able to maintain the products' competitiveness under the situation that the global competition involving the company's focusing segment of high display performance and high reliability expands and intensifies with significant falling in price.

EDT is developing a new series of display modules aimed as compatible replacements and offering the same features such as size and resolution and same mechanical outline structure. Until further notice they have planned for 3 sizes in XGA resolution (refer the table below), and more sizes can be added to this road map. If you are looking for other product sizes to replace Mitsubishi module, please do not hesitate to send us your inquiry and notify us what are the most important criteria for your application. Together we can develop the most suitable alternative.

Manhattan Skyline Ltd.
Unit 5, Bracknell Business Centre, Downmill Road,
Bracknell, Berkshire RG12 1QS, UK

www.mansky.co.uk
Tel: 01344 307733

SMARTPHONES just get smarter

With the roll-out of artificial intelligence accelerating and leading to the creation of new applications, so the market for smartphones continues to thrive. By **Neil Tyler**

Today's smartphones continue to push the envelope in terms of capability and new models are able to deliver not just high-quality photos and 8K videos, but can now support console-quality games and allow users to connect to the internet at multi-gigabit data rates.

The changed circumstances of the past twelve months has not only forced businesses to adjust to a new digital environment but has seen technological trends and innovation accelerate – whether that's the roll-out of 5G, the development of faster charging technologies, immersion technologies or new screens. All are having an impact on smartphones.

From new chip technology, such as Qualcomm's Snapdragon 888 chips and developments at technology giants like Apple and Google, where both are looking to develop and take control of their own hardware, to different styles of phone and the roll-out of cheaper handset designs the sector is witness to changing consumer requirements and new technological trends.

One of the most significant developments is that in 2021 most premium and mid-level smartphones will now come with 5G support.

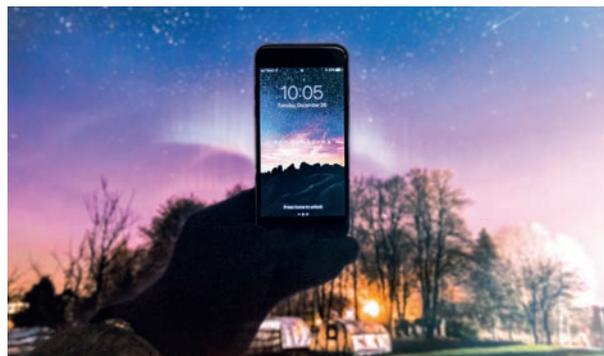
At the end of 2020, MediaTek unveiled the Dimensity 1200 and Dimensity 1100 5G smartphone chipsets - both of which are capable of delivering AI, camera and multimedia features for 5G enabled devices. These 6nm chipsets look to provide device makers with a variety of options when it comes to designing 5G smartphones and are able to support a broad range

of camera features, graphics and connectivity enhancements.

"MediaTek continues to expand its 5G portfolio with highly integrated solutions for a range of devices from the high-end to the mid-tier," said JC Hsu, Corporate Vice President and General Manager of MediaTek's Wireless Communications Business Unit. He continued, "The Dimensity 1200 is able to deliver 200MP camera support and advanced AI capabilities, in addition to innovative connectivity, display, audio and gaming enhancements."

The Dimensity 1200, for example, comes with an octa-core CPU that's been designed with an ultra-core Arm Cortex-A78 clocked up to 3GHz and, as such, can deliver much greater levels of performance.

Likewise, Qualcomm's Snapdragon 888 comes with integrated 5G on the SoC and the 5G modem no longer comes separately, as is the case with the Snapdragon 865 plus. This integration means that it will now be possible to develop more affordable 5G enabled smartphones – the improved computational capabilities



have removed the need to employ ever bigger and more sophisticated camera sensors.

The possibilities that 5G brings with it, such as better quality video creation and consumption, faster downloads, immersive AR/VR and gaming experiences, are considerable. Consumers will benefit from much faster services whether that's downloading content or sharing experiences of social media.

Both Google and Apple have moved into the development of hardware and in April last year Google unveiled a new project codenamed "Whitechapel." An in-house SoC, it is to be manufactured in conjunction with Samsung and follows the lead taken by Apple in developing and taking over the control of hardware.

Google's SoC comes with a number of cores that are dedicated to machine learning and the Google Assistant. Much like Apple the move reflects the company's desire to operate its devices under one ecosystem and, simply put, it wants better chips to deliver improved camera performance and overall usability.

Google's move mirrors that taken

or behavioural patterns of humans and can be used to solve problems that traditional computing has had difficulty in doing in the past.

The Huawei Mate 10 was the first mobile device to feature AI and today's smartphones using AI are better able to manage device performance and respond more accurately to their users' needs.

The impact of AI on the smartphone market has already been significant with Amazon's Alexa, or Apple's Siri AI-based voice assistants. By interacting with devices in other ways than through a graphical interface it is now possible for new applications to take shape.

"One of the many problems AI can solve is to make intelligent conclusions from images or videos such as facial recognition, object detection and tracking, motion detection and motion vector analysis," explained Johan Svensson, CTO of Imint, a Swedish company that specialises in intelligent sensor and data analysis.

"At Imint, we believe that with these building blocks there are even more possibilities for new applications within smartphones. We have already seen several applications such as image segmentation, face beautification, in addition to VR and AR applications like furnishing your home with virtual furniture, or trying on new clothes with virtual clothing.

"AI can be used to enhance traditional applications. This is especially true with the explosion in AI-processing power in new chipsets. That said, solving problems usually requires enormous amounts of data so that algorithms can learn to perform well, which puts certain companies, such as Google, in a strong position."

According to Svensson, AI can be used to address many different areas and applications.

"The specific application that pertains to Imint – and the best way to describe it – is that it's an extremely powerful technology that enables



"AI can be used to enhance traditional applications. This is especially true with the explosion in AI-processing power in new chipsets"
Johan Svensson

improved real-time video enhancement processing on a device in a way that anticipates the user's videography desires and/ or intent – and then auto-applies the appropriate adjustments to produce improved content.

"We've already been integrating AI into Vidhance - our flagship product suite. The Vidhance Selfie Mode combines AI and facial recognition to help ensure a user stays in-frame when video chatting or recording Instagram stories, for example."

By integrating AI and ML directly into smartphone software Svensson said it was possible to create more compelling video content and enjoy richer communication with friends, family, and co-workers.

"The ideal scenario for an end-user, when it comes to making a video, is when the technology knows what the user wants beforehand, and makes the requisite adjustments and corrections," Svensson suggested.

This could include technical adjustments, like auto-correcting a horizon or the auto-removal of noise and discolorations, or delivering creative adjustments, like auto-framing video subjects or auto-applying filters and effects.

"AI can also be used in up-scaling or super-resolution where users want to have many more details than what's actually physically captured by the sensor."

But that's only one of many benefits from using AI.

"There are a variety of benefits to the end-user. For one, the user will be able to have access to many new types of applications and features that will be able to gather even more types of information and intelligence."

The roll-out and availability of 5G will accelerate the use of AI, according to Svensson.

"With the bandwidth and low latency that 5G promises, we can start redefining the way in which network infrastructure is developed. Previously, with less bandwidth, many applications have been required to do calculations

by Apple which launched the A14 chip and M1 Macbook chip at the end of last year.

Artificial Intelligence

Artificial intelligence is now present across a growing number of smartphones. While it had traditionally been reserved for the most expensive devices it can now be found in more affordable handsets, delivering important new functions such as voice assistants, photography, object recognition and faster data processing. AI can be deployed to help users by learning their habits and tastes and then delivering suggested actions through different applications.

AI is able to mimic the intelligence



Left: One of the many problems AI can solve is to make intelligent conclusions from images or videos

in the edge, simply because it would take too long to pass the data from the edge to the cloud.

“With the speed of 5G, more of the computation can take place in the cloud, enabling more applications to send data to the cloud for processing.”

Consequently, with more data available more advanced applications that take advantage of data from more dispersed applications will be possible.

“More data is exactly what AI-applications hunger for, so we will see more advanced applications that take advantage of more data from more applications,” said Svensson.

When it comes to future applications, Imint is planning to incorporate AI into object recognition and object indexing technology. Though not yet on the market, the core technology has already been developed.

“We just need to find the right application and partner to bring it to market,” added Svensson.

Key trends

Among other key smartphone trends are likely to be an up-take in the availability of gaming and video applications.

The Coronavirus pandemic has certainly accelerated that trend whereby there has been a shift from console to mobile gaming and this is expected to continue over the coming year with the demand for more powerful smartphones with gaming capabilities expected to grow.

These types of smartphones will need to provide high refresh rate screens, use more powerful processors, offer ultra-fast storage and cooling capabilities.

The Snapdragon 888, for example, already comes with variable rate shading to improve gaming performance in smartphones and new smartphone models are expected to deliver more of these types of features.

But when it comes to gaming, for example, users also need improved



audio that is capable of delivering immersive sound and ultra-low latency connectivity. In response, Qualcomm has developed the aptX Adaptive audio technology which is able to provide a Bluetooth wireless solution that can work across gaming, video and music and has been designed to adjust performance based on the environment in which the user is using their smartphone and responds differently to what is being played on a device, whether that's video, music or gaming.

As Svensson explained, “Across the board AI has a significant role to play when it comes to improving the user experience and there are many ways to improve that experience. One specific use-case scenario that we strongly believe in is the ability for video to be utilized in a collaborative environment, where multiple producers provide their parts to a video. With AI engines helping to identify the relevant

Below: AI has a role to play in improving collaborative video experiences

parts of different producers' videos, we foresee that producers will be able to seamlessly merge these videos to create a professional looking end result.

“Imagine being at a football game where hundreds of people are filming from different angles. The ability to get a video from all different angles of the winning goal is something that many would want.”

All these improvement will require power and charging technologies are evolving with fast-charging products coming to market which can charge a smartphone from zero to 50% in just a matter of a few minutes.

While the technology is improving there is a call for OEMs to make this technology safer for the user, minimising any potential safety hazards and risks, while at the same time providing an efficient and secure charging experience.

A number of companies have developed battery safety monitoring chips that are capable of monitoring a mobile phone in real time to see if it has been damaged.

The smartphone market continues to evolve rapidly but as consumer demand continues to change, so smartphone brands will need to come up with newer and more innovative technologies to meet changing market requirements.





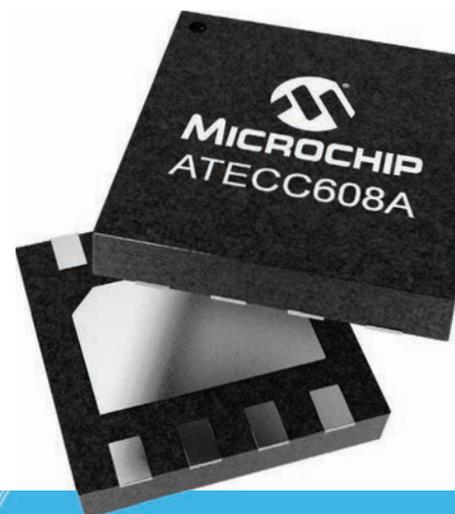
Pre-Provisioned Security Solutions Within Reach

Trust Platform for the CryptoAuthentication™ Family

The Trust Platform for the CryptoAuthentication™ Family enables you to integrate hardware-based secure key storage with pre-configured or preprovisioned devices for any low- mid- and large-sized device deployments. Three platform options allow you to choose exactly the authentication model that fits your system needs: Trust&GO is pre-provisioned with default certificates and keys; TrustFLEX is pre-configured with the most common authentication use cases and allows you to use your preferred certificate authority; and TrustCUSTOM offers you a fully customizable solution.

Key Features

- Trust&GO - pre-provisioned
- TrustFLEX - pre-configured with the most common use cases
- TrustCUSTOM allow you to dial in your customization needs
- Common criteria JIL-rated high secure key storage



microchip.com/TrustPlatform

The Microchip name and logo and the Microchip logo are registered trademarks and CryptoAuthentication is a trademark of Microchip Technology Incorporated in the U.S.A. and other countries. All other trademarks are the property of their registered owners.
© 2020 Microchip Technology Inc. All rights reserved.
MEC2356A-ENG-12-20

Delivering the infrastructure for 5G

Brendan Farley, MD of Xilinx EMEA, discusses the phased roll-out of 5G infrastructure and the opportunities that will create. By Neil Tyler

Based in Ireland, Brendan Farley is managing director of Xilinx EMEA and vice-president for wireless engineering with responsibility for the development and roll-out of technologies that will support the deployment of 5G infrastructure.

Appointed to the position last year he has to manage teams and a complex R&D process that spans Ireland, the UK, Germany, US and India while ensuring that all the company's key functions, whether that's R&D, sales, marketing, HR, IT, finance and legal are working effectively together.

"Our work is focused on developing the infrastructure that's needed to deliver 5G and that means working very closely with our customers and partners, as they look to deliver new applications and services," Farley explains.

"Xilinx's products are used in cellular base stations, and the roll-out of 5G has proved a significant opportunity for the company and much of the R&D activity takes place here in Ireland."

Prior to his appointment Farley held a number of senior engineering management roles within Xilinx, and before joining the company had worked for many years in the Irish technology sector.

"It's been an interesting few months and my appointment, which coincided with changed working practices as a result of the pandemic, has meant that it's been even more important to stay in touch with the team and keep people motivated and informed.

"Likewise with our customer engagement we've probably been far more active using tools like Zoom, but I think there's now a real desire to get back to normal and re-establish those personal relationships that are so important in this business. However, we don't want to lose the benefits working online has brought, whether that's breaking down old formalities of lots of people having to turn up on site to make decisions or simply enabling decision-making to be faster."

Turning to 5G and the roll-out of the infrastructure necessary to support it, Farley says that the acceleration in the commercial deployment of 5G is driving the market forward.

"The commercial roll-out of 5G has certainly picked up and as a consequence we're seeing increased investment in network infrastructure in order to meet the needs of 5G users. Xilinx has been front and centre here, helping to upgrade networks across China, South Korea and in North America."

The strength of Xilinx's portfolio lies in the programmability and adaptability of its FPGAs and SoCs.

"Specifications are constantly changing," Farley points out, "along with the features that need to be supported. It's difficult to keep up with

these changes if you're designing with ASICs, so the flexibility that our adaptable SoCs and FPGAs provide has made them far more suitable."

A phased roll-out

According to Farley the roll out of 5G will need to be phased.

"The first phase is about upgrading the infrastructure from 4G to 5G in order to support traditional operator business models. Our technology is being used to increase capacity and reduce latency and that's happening now – and we're playing very strongly in that.

"For operators though that's not enough. They will not consider 5G a success until they are able to develop new revenue streams from new verticals – and that will require heavy investment. But that's essential and that's what we'll see in Phase 2."

Farley suggests that Phase 2 will see new revenues from new applications which will be driven by the operability and integration of various OpenRAN solutions.

"This whole process is closely tied in with OpenRAN, one of the cornerstones of which is the separation of hardware and software, particularly at the edge – in other words the base band next to the tower. At the moment the use of ASICs is widespread and there's a focus on proprietary hardware. We are now seeing a move to the virtualisation of the base band with software doing most of the signal processing that was traditionally done by an ASIC."



“The first phase is about upgrading the infrastructure from 4G to 5G in order to support traditional operator business models”

The advantage of OpenRAN and the use of software is that it is now possible to push third party applications on to servers which will be able to process data in real-time.

“That will open up a host of opportunities for new applications,” Farley suggests. “Just look at the automotive space and the concept of cellular vehicle to everything – if you have low latency all the traffic will be passing through a base station tower to another end user. That will make it possible to implement augmented reality, for example, in the form of holographic navigation systems. And, if you can connect cars to a 5G tower, you can connect traffic lights and pedestrians and if you are able to deliver latencies of under 50ms, it’ll be possible to significantly improve road safety.”

There is, however, still some way to go before we are able to deliver these new services. Many would argue that we need to be careful about not over-promising what is deliverable and being realistic about possible timelines, especially when it comes to 5G standalone services.

“For sure, I would agree with that,” says Farley. “In terms of a basic route map we need to upgrade the infrastructure first and then virtualise the base band and that has started to happen. Samsung has already demonstrated that it’s possible to get the same level of performance from a virtualised system, but not everything can be run in software. Some functionality is only possible through hardware.

“This is a vision that will be captured in new 3GPP standards as they emerge but we do need to be able to demonstrate increased 3X capacity with the existing business models; then we need to virtualise and disaggregate the network, especially the base band with some acceleration from FPGAs, and then we need to engage with the verticals, whether that’s factory automation or autonomous vehicles.

“We need to get buy-in from those new verticals and moving to ‘server based architectures’ will be critical. For factories that could result in a move to a wire-less factory. That will require acceleration at the edge, but once the software platform is in place people will start to see the opportunities.”

A successful deployment, however, will depend on an ability to provide low latency especially, for example, in a factory environment that’s looking to operate automated processes, robots and machine vision applications.

“All of this requires the 3GPP to reach out to these verticals, to see where new value can be created and what the new opportunities are for operators,” says Farley. He continues, “The key thing is that all of this remains a moving target.

“Whether we are talking about standards or features, something is being added every year and that’s why the programmability and adaptability of the devices that Xilinx provides are critical when it comes to 5G.”

Finding that 'unexpected' bug

It's a real challenge to identify bugs in embedded code across the Internet of Things, but visual trace diagnostics has made that process easier. By **Johan Kraft**

As more projects move to the Internet of Things, the challenges of embedded code become more obvious. With millions of devices out in the field in all kinds of conditions, bugs in the embedded code are even more likely to show up at one time or another. It is also more difficult to identify and fix those bugs when the embedded systems are out on the road or part of a supply chain on the other side of the world.

This is potentially a huge problem. Research from 13,500 software projects found that the defect removal rate was 95%. While this sounds pretty impressive from the engineering perspective, it still means 5% of defects end up out in the field. With 50-100 bugs typically introduced per 1000 lines of code, this means there can be 25 to 50 major bugs in a system of 50,000 lines of code.

As projects get bigger, combining embedded control and cloud systems across the latest industrial automation systems and the IoT, this problem keeps getting bigger, too. While engineers do their best not to introduce bugs in the first place and use extensive testing to remove all the bugs they can find. And this is where the key lies – how to find the bugs, especially when they can show up unexpectedly.

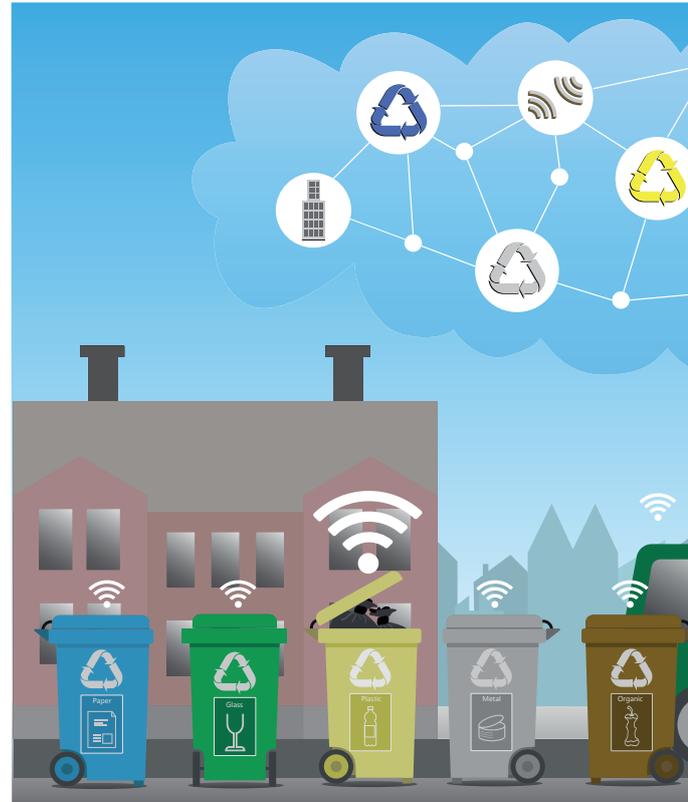
When embedded systems are deployed in things that move around, whether that's in smart bins for collecting refuse or monitoring the cold chain delivering vaccines to tackle the Covid-19 pandemic, all kinds of different situations may occur. Some

things can create a problem, and the developers behind the software need to know.

The bug may not be obvious, but it may be a precursor to a bigger problem. Catching problems early before the issue becomes catastrophic is ideal, especially if this can be done without the customer or user even noticing.

The same forces acting to drive the connectivity of embedded systems can also be used for identifying and reporting any potential problems. For example, monitoring the ambient and air temperature across an entire cold chain distribution highlights weak points that can be addressed, and ensures that products, including vaccines, arrive effective at the point of use without having been compromised. This requires extensive monitoring and logging, with data uploaded to the cloud for analysis.

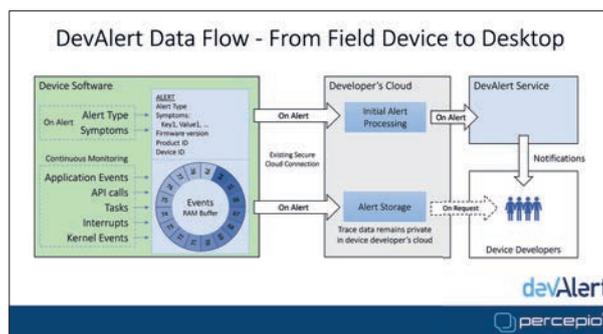
Similarly, smart refuse systems use sensors and wireless links to determine when a bin is full and needs to be collected, rather than relying on a set schedule. Fleet management software and intelligent routing then determine the most cost-effective



retrieval method, day by day.

The cost of a bug in such systems can be dramatic. While one trash bin not being collected is not a disaster, angry calls from customers with overflowing bins is not what an operator wants to deal with. It is also vital to know whether a sensor node has failed as a result of the sensor, the link or other parts of the system.

Losing a batch of vaccines that might have been compromised during transit is much more significant, especially if this is the result of a sensor failure. And if a sensor fails, there is no way of knowing whether the batch has been compromised or not.



Catching errors

So catching errors as they occur and notifying developers is of tremendous value to enterprises. The average cost of fixing bugs in an embedded device's first year of service alone can run into hundreds of thousands of dollars, and this does not include damages caused by bugs.

One method that is of tremendous



Solveig/stock.adobe.com

value to the developer at the benchtop is visual trace diagnostics. This gives a visual timeline of the internal software events during operation and is a key tool for a developer tracking down bugs. The trick is to bring this data out during deployed operation and provide the developer team with detailed information about software issues in the field.

This is where the cloud integration comes in. As more nodes are connected in the IoT, that same channel can also be used for diagnostic data. The cloud connection can be used to alert developers when an error is first detected and provide visual trace diagnostics to identify the root cause.

This is a powerful concept that not only highlights when a node fails, but it also provides an explanation – the timeline of software events just before the issue was detected, showing what led up to the problem. This way, the software can be corrected and updated quickly, while only a few customers have been affected by the

bug, and thereby avoid a much larger problem.

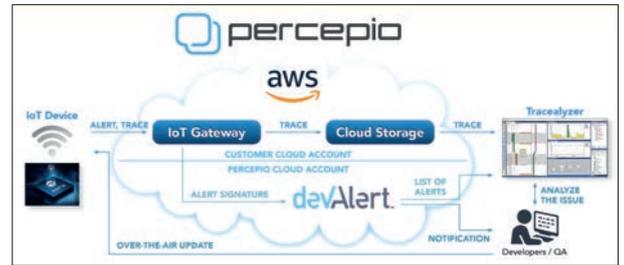
This concept is the heart of a tool called DevAlert. In combination with the visual trace diagnostics technology Tracealyzer and a sophisticated cloud data management system, it can be used to monitor any kind of IoT connected device.

When something unexpected happens, for example if the system automatically rebooted to recover from an error, the events that led to the reboot are captured and available. But it is the awareness that is the key. When developers receive an alert, they can immediately look at the diagnostic trace information and see exactly what happened.

By including a trace with the alert, it becomes much easier to identify the situation that caused the problem. When you add visual trace diagnostics, including many types of visual overviews, it becomes even easier for developers to understand the problem. This allows developers to find the root cause and fix the problem quickly. The reaction time matters. Most bugs in deployment don't show up directly for all users, otherwise they would surely have been found during the testing. So the faster an update can be provided, the fewer customers will be affected.

This is particularly helpful as the customers won't necessarily report a problem and provide the information that a developer needs to reproduce and fix the bugs, especially for consumer devices. Even when a bug is reported, it can be with very vague information, leaving a developer at a loss on where to start.

The concept involves three software components. The DevAlert Firmware Monitor (DFM) is a compact software library that device developers embed in their RTOS-based IoT application. This agent keeps a trace of recent software events and provides a way for error-handling code in the application to report any condition of relevance to the device



Above: DevAlert can be used to monitor any kind of IoT connected device

developer, errors, proactive warnings or other diagnostic information, whether related to software or hardware. The alert message is then uploaded to the cloud account of the device using an existing secure connection, such as MQTT over Transport Layer Security (TLS).

But the DFM monitor on its own could just deliver a deluge of data. The DevAlert cloud service takes that data and looks at error codes and any other symptoms and notifies the developers in case of a new unique issue if a new combination of symptoms has happened. This avoids the problems of duplicate alerts flooding a developer. Percepio's Tracealyzer tool can then be used by the developer back in the lab to analyse the provided trace.

It is important to understand that DevAlert in no way replaces conventional testing; you need both, just as most cars have both seat belts and air bags. Good, systematic testing typically removes 95% of the bugs, as we stated above. Error reporting in the field will help you catch those bugs that testing couldn't find, the difficult ones that only appear under certain conditions. There is often an astronomical number of potential scenarios in the software, that depends on the inputs, the software timing, device settings and other environmental factors (e.g. the Wi-Fi connection). Any one of these may have latent bugs, that may cause the device to crash or produce incorrect data. For example, some device might fail if the Wi-Fi connection receives multiple packets within 5 milliseconds while it is busy writing data to a flash memory.



Author details:
Johan Kraft is CEO and Founder at Percepio



Cooking up a connected product

Connected devices tend to sit at the more complex end of the embedded system spectrum, irrespective of whether they gather sensor data by themselves or interface with pre-existing equipment.

Their complexity arises from the interaction between their three constituent parts: the local part, which generates the data or needs some control; the far system which records the data and/or provides instructions; and the connection between them, which typically operates over a large distance and is often observable (and even interruptible) by untrustworthy third parties.

These complexities are, however, often outweighed by the improved customer engagement, detailed business insights, and even entirely new commercial models that can be achieved.

The fine balance between complexity and function means that building the right kind of connected device is often part art, part science. In fact, it's a lot like cooking a great meal – there are practically an infinite number of ways to bring key ingredients and dishes together. Some of which can be assembled quite quickly, whereas others will take more time and effort to perfect, but which may be better suited to someone's tastes.

One approach is to break the IoT development process into several

Adding a connected device to your portfolio can be challenging. **Jonathan Pallant** looks at how to pick the 'perfect ingredients'

specific but fundamentally connected choices. But it isn't always easy to work out which decisions to take first, and there is always a risk that an early decision could restrict your later choices and lead to the whole system being sub-optimal.

We've alliteratively grouped these choices into six different areas to consider when planning the perfect IoT 'feast'.

Connectivity is probably the best place to start, as it's the fundamental component that the rest of the offering hangs around. The main course, if you will.

The Connectivity Course

The key connectivity technology requirements for products are defined by cost, power consumption, range, data rate, and latency. But they are ultimately tied together by the Laws of Physics and Economics. And all of today's technologies exist in their own particular sweet-spot within this multi-dimensional space.

Most systems will probably need wireless connectivity to access the Internet, which sadly is typically less reliable and higher cost than using a wired connection. Almost all wireless standards involve at least two classes of device – a smaller, lower-power component

(like a mobile phone) and a larger, higher-power gateway (like a Wi-Fi router). You also need to consider the ownership and reliability of those gateways, comparing for example a homeowner's Wi-Fi router versus accessing a mobile network.

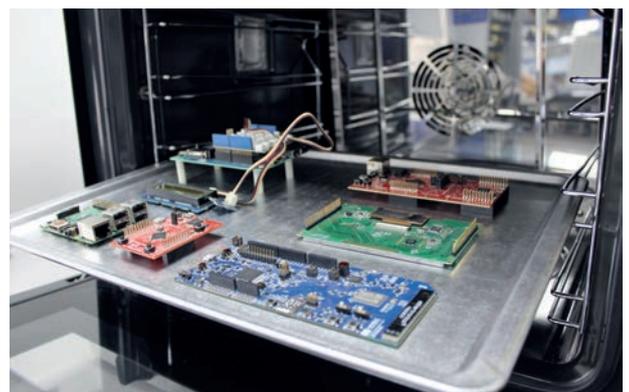
Whichever way you go, it is almost always better to rely on existing connectivity standards with broad market acceptance rather than developing something from scratch.

The Chipset Course

The first decision to make is whether to opt for a chipset that integrates processor, memory and radio, or whether to combine several ingredients to get the right mix of performance, price and power consumption.

If you need low-power cellular

Below: There are few absolutes in technology. There will always be a range of potential solutions available



connectivity, either with LTE-M or NB-IoT, a device like the Nordic nRF9160 is a compelling package of CPU, RAM, Flash, GPS and Modem. But, if you can afford the PCB space, you might prefer the flexibility of a “standard” microcontroller and a separate LTE module. A similar story applies to wireless standards like Wi-Fi and Bluetooth, where there are microcontrollers available with the radio built-in, and also stand-alone modem modules.

Of course, IoT devices don’t just run on basic microcontrollers. There are plenty of examples where the power and technical ability of a full Linux Kernel-based system was the right choice, even with the knock-on effects on power consumption, space and software support burden.

The Core Course

This stage involves bringing in pre-existing software (essentially “third-party IP”) including: the Operating System (or Real-Time Operating System) kernel; the chipset drivers; and the board support package (BSP) which customises those drivers for your particular PCB design.

Chipset vendors often provide a free Software Development Kit (SDK) containing all of this. But it is worth remembering that this SDK may not work with anyone else’s silicon, limiting your ability to change supplier in the future. If that matters, perhaps consider a vendor-neutral offering such as Amazon FreeRTOS or Microsoft’s ThreadX.

Linux-based systems have similar options too – often with a free “distribution” to get you started. But you may prefer to roll-your-own or to use a more generic off-the-shelf Linux distribution, such as Ubuntu Smart Start or Fedora IoT. It all depends on your security requirements, development timescales and any system limitations on power consumption and storage space.

The Code Course

Of course, you can’t just install embedded Linux on your device, or flash on an RTOS, and call it a day. You’ll need to add the “secret sauce” by writing code to customise this system to your needs. That could be as simple as specifying the endpoints and keys to upload sensor readings, or as complex as running multiple bespoke protocol stacks and sophisticated AI processing at the edge for good measure. Either way, the language and tools you use will have a huge impact on your up-front development programme, as well as for on-going support and maintenance.

The C programming language has been the default for almost 40 years now, but maybe there’s value in trying a memory-safe systems language, like Rust or even something like MicroPython. But if you find your Chipset or Core can’t handle your preferred choice here, you may need to go back and re-think.

The Cloud Course

Most connected device developments involve the gathering of data for later analysis, and the issuing of commands to control devices “in the field”. It makes sense almost universally to offload this computing requirement to one of the pre-built IoT management and data storage offerings available from the big “cloud” providers, rather than using on-premises systems.



Author details:
Jonathan Pallant is senior embedded systems consultant at 42 Technology

Whichever cloud platform you choose though, you’ll need to balance the up-front costs of integration against on-going support and maintenance. And remember, you’re going to need support from your cloud, chipset and connectivity providers for the entire lifetime of your product in the field.

And finally...the Comms Course

The default Communications protocol your Code will use to talk to your Cloud provider will usually be text-based (such as JSON or XML) running over an encrypted connection-oriented HTTP link.

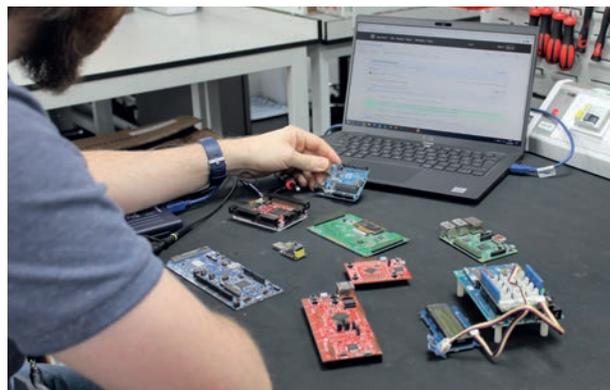
Whilst this sort of stack is ubiquitous (it’s how most current websites are built), its plain-text nature increases both bandwidth requirements on your connectivity and chipset resources, and the connection-oriented nature can interact badly with some high-latency wireless services. For example, if you are designing an ultra-low power monitoring system using NB-IoT, you might be better with a binary connection-less protocol like CoAP.

Also, being able to rely on a well-used, real-world-tested protocol stack (such as the one that comes with your Core software) is almost certainly going to be better than trying to spin one yourself.

It’s worth remembering

There are few absolutes in technology and there will always be a range of solutions to meet your requirements. What works best will depend on your budgets, timescales and the skills you have available in-house, from an external development partner or through any existing relationships with hardware, software or network providers. But, even the best technical solution will never become a winner unless it gets to market on time and with the right combination of price, specification, and availability too.

Below: Deciding what to prioritise is not always easy when developing new connected products



Addressing safety critical FPGA designs

Advanced EDA tools and methodology is helping designers to avoid synthesis bugs in safety-critical FPGA designs, as

Rob van Blommestein explains

Field-programmable gate arrays (FPGAs) are the dominant hardware platform in many safety-critical, low-volume applications, including aerospace and nuclear power plants (NPPs).

Modern FPGA devices feature integrated microprocessor cores, digital signal processing (DSP) units, memory blocks and other specialised intellectual properties (IPs) and these advanced devices allow for the implementation of large, high-performance system-on-chip (SoC) designs with integrated safety mechanisms, making a strong case for adoption in additional safety-critical applications traditionally dominated by application-specific integrated circuits (ASICs).

From a high-level perspective, the FPGA and ASIC development flows are similar. Register-transfer level (RTL) coding and integration of third-party intellectual properties (IPs) are crucial steps in the front-end part of the flow. Extensive functional verification of the RTL design model reduces the risk of mismatches between requirements and RTL behaviour. At this stage, specification, coding, and module integration mistakes are the main source of systematic faults that, if undetected, could lead to dangerous failures of the FPGA device in the field.

The RTL model then goes through several implementation steps (see Figure.1). Synthesis and place-and-route tools map the design onto the target FPGA device. The bitstream generation step produces the file used to program the FPGA. Functional

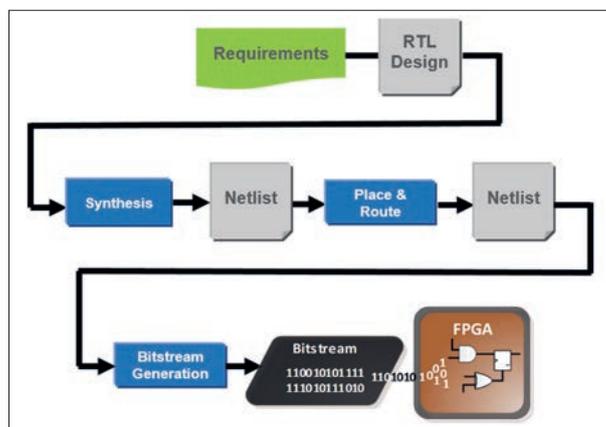


Figure 1: The FPGA development flow

verification of implementation steps reduces the risk of mismatches between the derived netlists and the RTL design model. This is crucial to close the loop and ensure that the functionality implemented in the FPGA device matches the hardware requirements.

Functional bugs can be introduced during implementation steps, either because of RTL issues that cannot be detected during synthesis, or because of malfunctions in implementation tools, particularly synthesis and place-and-route, corrupting the original RTL functionality. While engineers may expect that implementation tools have been extensively tested prior to release, each design and coding style are unique and may trigger unknown corner cases.

Regardless of the implementation tool used, there are coding issues that may go undetected during RTL verification and creep into the synthesis netlist. In some corner case scenarios, RTL simulation behaviour may not match the

behaviour of the corresponding netlist in the presence of unknown, or X, values. Consequently, while the synthesis tool operates correctly, its generated netlist may still not match the intended RTL behaviour.

Requirement-based testing does not explicitly target these type of corner cases, which can therefore be missed. RTL linting tools may warn about these scenarios. However, they provide neither a definite answer nor a simulation trace that shows how the bad scenarios may occur.

Desired functionality

FPGA implementation tools must fit the desired functionality into a prefabricated structure while meeting performance and power consumption goals. Implementation tools perform significant changes to the original logic structure of the design to improve device utilization and overall quality of results (QoR). Advanced optimisations pose a higher risk of corrupting the RTL functionality.

Certain FPGA synthesis tools support the automatic insertion of hardware safety mechanisms. Safety mechanisms shall not change the design functionality when no fault is present. In the event of a random hardware fault occurring during field operation, they need to raise an alarm and potentially correct the effects of the fault on the fly.

Synthesis tools may transform the encoding of finite state machines (FSMs), for example to include Hamming-based error detection and correction. Triple modular redundancy (TMR) is another type of safety mechanism where a critical logic function is triplicated and voting logic added to determine which of the three outputs should be considered as correct. Logic duplication and inference of memories with error correcting codes (ECC) may also be supported. Users certainly benefit significantly from these design enhancements that can be performed automatically by the synthesis tool.

Author details:

Rob van Blommestein, Vice President of Marketing, OneSpin Solutions.

However, the inserted logic could contain bugs and must be rigorously verified.

An effective verification flow must detect bugs as soon as possible once they are introduced. Detecting an RTL issue or synthesis bug during lab testing or Gate Level Simulation (GLS) of the place-and-route netlist is inefficient. Moreover, the effect of bugs introduced by implementation tools is unpredictable. Simulation tests are not intended to verify the correctness of the implementation tools, and even running all available tests at gate-level only provides limited confidence. This approach is indeed far from exhaustive. Finally, debugging GLS and lab test failures is hard and time consuming.

Formal methods are widely recognised as a powerful and exhaustive verification technology. Formal design inspection and exploration is valued for detecting both basic and corner-case RTL issues early and without the need for a simulation testbench.

Automated, formal inspection of RTL code detects issues before synthesis starts. Unlike linting, formal tools provide a definite answer on whether an array may be indexed out of bounds. In this case, the tool provides an easy-to-debug simulation-like trace, or counterexample, that demonstrates how the design misbehaves.

EC tools

Formal EC tools can mathematically prove (or disprove) that two designs are functionally equivalent. This is the most rigorous way to ensure that synthesis and other implementation steps have not introduced bugs. The input design to the implementation tool is typically named golden design. The generated netlist is named revised design.

Combinational EC largely relies on one-to-one mapping of states between golden (e.g., the RTL) and revised (e.g., post-synthesis netlist) designs. Through state mapping, the complex

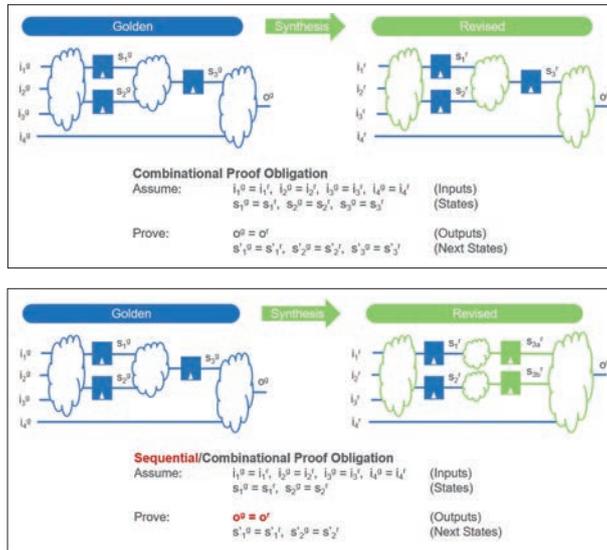


Figure 2: Proof obligation in sequential and combinational equivalence checking

problem of proving that two large designs are functionally equivalent can be split into a multitude of much simpler problems: comparing the functionality of two combinational logic cones (see Figure 2).

The design transformations performed by FPGA implementation tools significantly break one-to-one state mapping. Formal sequential EC algorithms can prove equivalence of sequential logic cones, thus not requiring state mapping. However, while these algorithms have improved dramatically in recent years, they do not scale. Partial state mapping is necessary to leverage combinational EC wherever possible and apply sequential algorithms only on limited design portions (see above). In this context, identifying corresponding states is a crucial, challenging task. Manual mapping is tedious and time-consuming. Mistakes waste engineering resources.

Formal verification signoff (see

Figure 3) enables engineers to use advanced FPGA optimisations and the latest synthesis technology with confidence. Formal RTL design inspection is more powerful than linting and finds issues early, prior to synthesis. Formal EC proves that the golden design functionality is not corrupted by the implementation step. Finally, formal fault injection and verification supported by specialised safety tools, can automate the verification of safety mechanisms in the scenarios when faults occur. With this flow, weeks of GLS and lab testing can be replaced with hours of formal tool runtime.

FPGAs have long been the hardware platform of choice in many low-volume safety-critical applications. Nowadays, these devices can implement complex functions while fulfilling tough performance and power goals, competing with ASICs also on high-volume safety-critical applications, including automotive.

The availability of advanced EDA tools and methodology is crucial to support this trend. ASIC development has used formal EC for nearly 20 years. Automated formal checks prior to synthesis are also widely adopted by ASIC teams. The same technology is now available in FPGA development, enabling a robust, efficient implementation process. OneSpin's formal signoff flow of FPGA implementation has been designed to be orders of magnitude more rigorous and efficient than GLS and lab tests. The technology is mature and proven on hundreds of industrial designs for communications, NPPs, and other safety-critical applications.

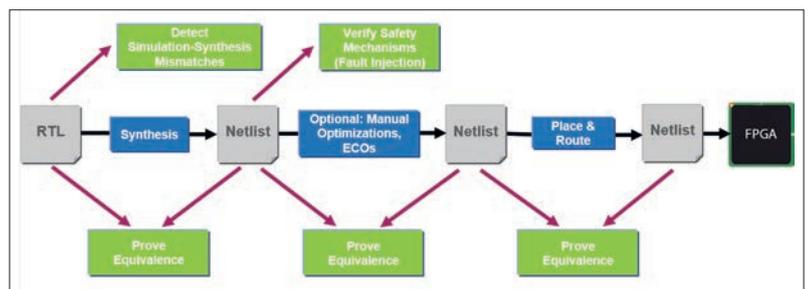


Figure 3: Functional sign-off of FPGA implementation with formal verification

Thermal evolution

Generative design and machine learning point to novel ways to cool systems. By **Chris Edwards**

Two years ago, a team of HP was faced with the challenge of getting more airflow out of a duct used to cool the print-heads of an upcoming inkjet printer. Without efficient cooling to keep the ink flowing at the right temperature the printer cannot operate at full speed. Simulations showed the traditional duct design was causing too much recirculation and turbulence to work effectively.

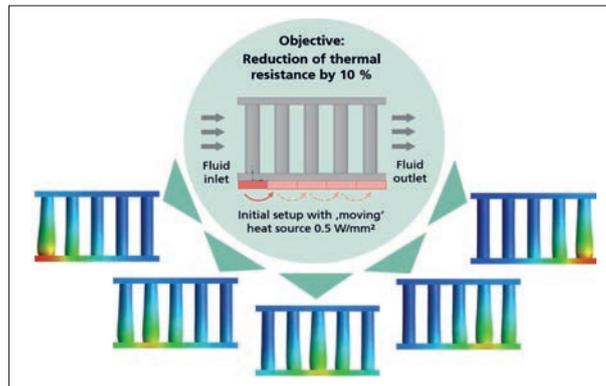
To get around the issues, HP worked with engineers at Siemens Digital Industries who coupled fluid-flow simulations with a technique mechanical CAD companies now term generative design to come up with a radically different structure.

“As the design was shaping up it kind of defied whatever everyone included me expected,” said Siemens simulation engineer Julian Gaenz after the launch of the printer that would use the duct, pointing to a tongue-like protrusion in the final design that its predecessor lacked.

There are two main tactics used in this kind of generative design. One is to randomly generate variants after each simulation run in the expectation that one or more will work better. Another is to home in on problem areas, such as the recirculation in the HP printer’s duct, and come up with shapes that seem likely to reduce the problem before simulating a complete part to see how well the changes worked. HP is far from alone in exploiting simulation-driven design.

Tom Gregory, product manager of thermal-analysis tools supplier Future Facilities says his company’s 6SigmaET tool has been used to help determine the optimum shape for pins used in that style of heatsink.

“The optimisation was performed by



using genetic algorithms in combination with CFD. The aim of the optimisation was to reduce the pressure drop across the heat sink while maintaining thermal performance.”

Though it is entirely possible to create the variants or iterations by hand, machine learning looks to be a useful tool for generative design because, in principle, it can create variations quickly and come up with counterintuitive but effective unexpected results. Gregory points to several research papers have come up with novel heat-sink designs.

“The designs that are generated are impressive and visually stunning. However, they cannot be manufactured by traditional methods.”

3D printing not essential

In terms of research heatsinks often need 3D printing, as did the HP nozzle. However as that was replacing six injection-moulded parts, the newer part worked out more cost effective once the 3D printing process was optimised.

Lieven Vervecken, CEO and co-founder of Belgium-based generative-design startup Diabatix says 3D printing is not essential.

“We have a few manufacturing techniques that we support by default,

Above and below: Work by the Fraunhofer IISB to determine the optimum shape for pins used in heatsinks

such as CNC machining, extrusion, die-casting and sheet metal forming,” Vervecken says. Each has its use in different markets. For example, sheet metal forming has proved useful for making the large battery cold plates needed by electric vehicles.

Diabatix has built a generative environment around a selection of open-source design and simulation tools that iteratively refines the structure of a heatsink or cooling component based on the toolsuite’s analysis of thermal flows and mechanical structure.

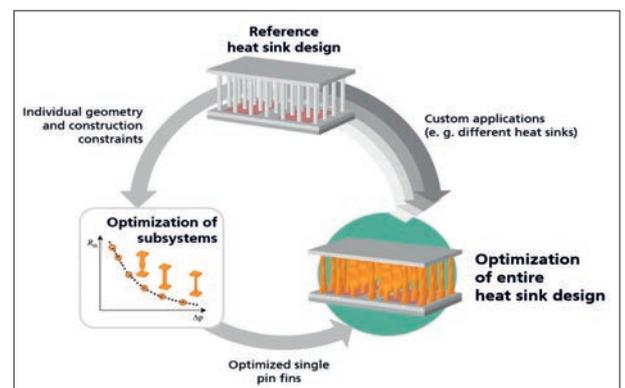
The company is launching a cloud-hosted version of the tools in April.

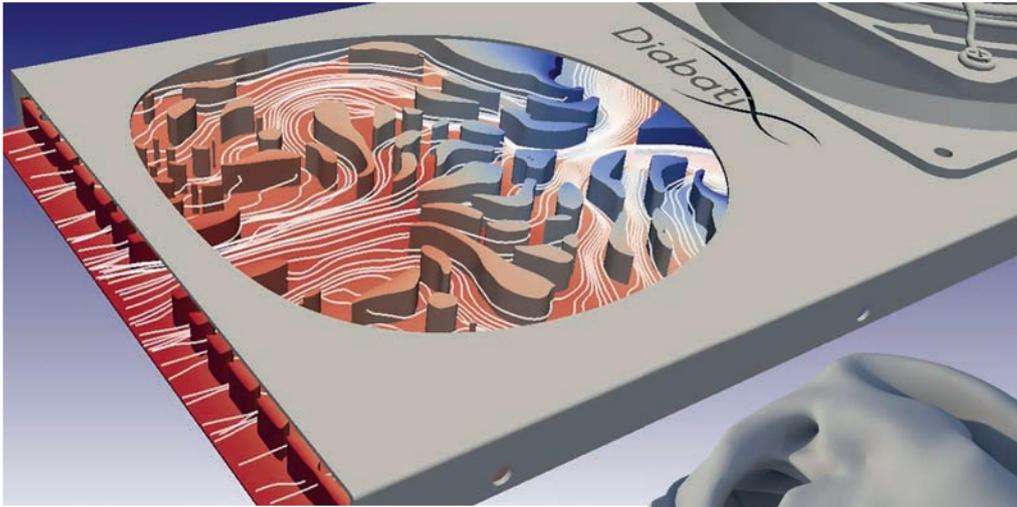
Although generative design can use randomisation to create novel designs, this was not the approach taken by Diabatix. The problem with randomisation is that the simulation overhead needed to assess each variant makes the technique too unwieldy. Instead, Vervecken says the approach used by the start-up uses machine learning to build algorithms that can come up with effective but often unexpected shapes and structures rather than the parallel fins found in most conventional designs, with manufacturing constraints used to stop them from becoming too impractical to manufacture.

“This is why we can generate up to 30 per cent additional cooling performance,” Vervecken claims.

Thermal engineering

Another potential target for machine learning in thermal engineering is to reduce the bottleneck introduced by





simulation if a large number of variants need to be analysed in parallel before further progress can be made.

Training a deep-learning model with the results of extensive simulations has been used in areas such as materials and particle physics research because the trained models can run faster than the original simulations, which may take days to complete even on a supercomputer. They do not offer full accuracy but can be used on a fast path to home in on parameters before running a final detailed simulation.

GPU maker nVidia, which is already a strong proponent of machine learning has applied the same technique to a project called SimNet to do the same for the kinds of computational fluid dynamics (CFD) models used for thermal engineering.

A key problem is that deep learning itself is computationally intensive and requires weeks of simulation to build a big enough training set to let an AI model learn how parameters affect air and heat flow in space.

Vervecken says the range of projects in which they are involved, which range in scale from mobile-phone processor heatsinks to truck-scale battery cold plate, increases the complexity of the problem.

"To train a single model to cover that whole range? That's difficult," Vervecken says, saying that it makes more sense to prioritise machine

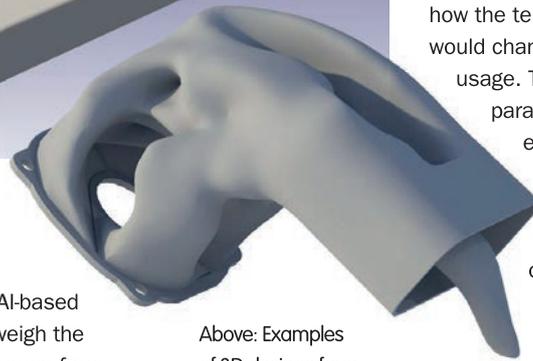
learning for the design tools rather than the simulation engine.

Against a possible AI-based speedup you need to weigh the contributions that can come from simply optimising the computational fluid dynamics (CFD) algorithms directly.

Specialist vendors such as Future Facilities have employed multicore execution and other optimisations to cut overall turnaround time. In a study with Rohde & Schwarz, the time taken from run a simulation from CAD import through mesh generation to analysis was cut to below 15 hours from more than 40.

As well as heatsinks, simulation-driven machine learning can be applied to core device and system design.

"One area that Future Facilities is working with machine learning research groups is in using CFD simulation to



Above: Examples of 3D designs from Diabatix

train machine learning algorithms," Gregory notes.

There are a number of areas where the machine learning can prove useful. One is to help build simple models of heat generation and transfer that can guide algorithms used to decide when to power down processors on a multicore SoC.

"Thermal simulation can easily train the machine learning algorithm how the temperature of the device would change based on processor usage. This can be done quickly in parallel before the device has even been manufactured," Gregory says.

Another use lies in R&D for new types of device, such as those used for power electronics, where excessive heat can cause thermal

runaway and other safety issues. "Comprehensive training data is essential for an effective machine-learning algorithm, but it's often not practical or safe to obtain training data from physical devices. If you are optimising a new device or system, a physical device may not have been created," Gregory says.

Performance is not the only motivation for using AI-based models. "Very often it's a cheaper solution they are looking for," adds Vervecken. "There's a misconception that exists with our technology. We are not always looking for the most efficient heatsink. We have a project on-going where the customer is instead looking for the most affordable solution. The customer said performance is not an issue: the existing design is already overprovisioned by 30 per cent but they can't get that type of design to be cheaper."

It is still early days for machine learning for thermal engineering but generative design techniques are demonstrating the value of simulation and rapid iteration in coming up with designs that work better, work out cheaper or both.



Below: The result of generative design – the HP duct tongue profusion

START GROWING YOUR BUSINESS NOW

For over 35 years, MA Knowledge has been supplying the engineering and manufacturing industry with high quality, targeted business leads

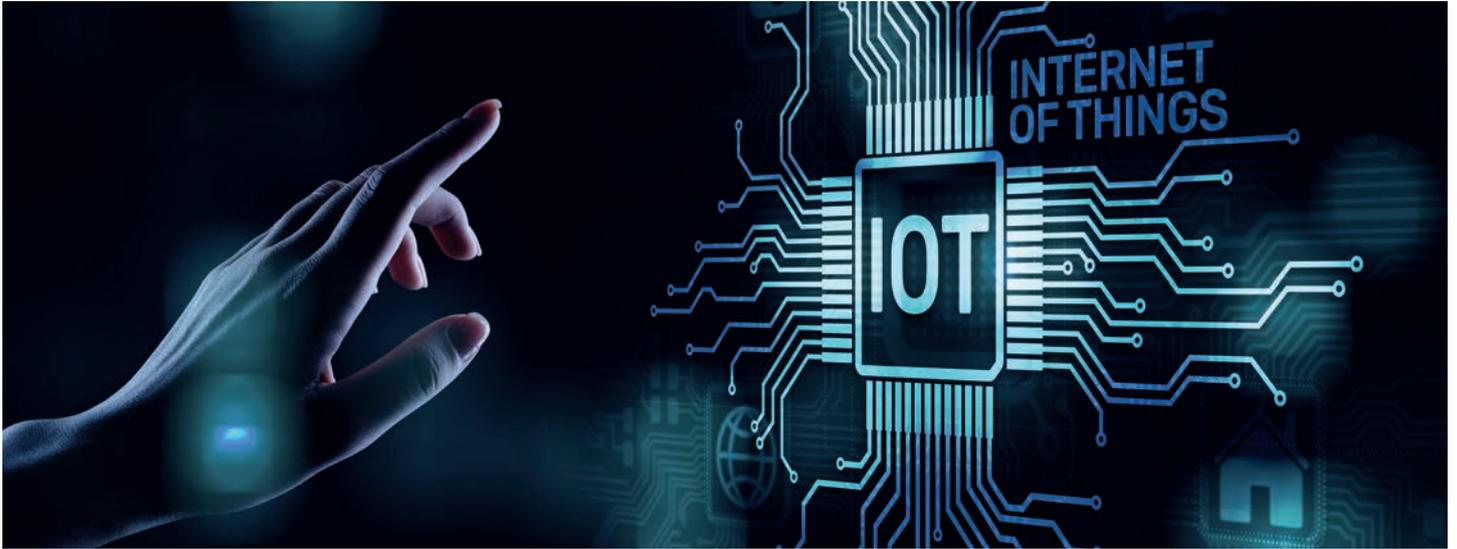
- ✓ Data rental
- ✓ Data multi use licences
- ✓ Email marketing
- ✓ Direct mail
- ✓ Telemarketing
- ✓ Lead generation
- ✓ Sales prospecting
- ✓ Data audits/ cleaning



PROVIDING INTELLIGENCE LEAD MARKETING SOLUTIONS

We work in partnership with you to deliver high value, accurate data which is tailored to your requirements to find out how we can help improve the value and power of your marketing messages contact: **James Creber** on **01322 221144** or **james.creber@markallengroup.com**

www.ma-knowledge.com



A BLENDED STRATEGY

Why using a blend of networks will deliver a more successful IoT deployment. By **Nick Sacke**

IoT production deployments are underpinned by a carefully orchestrated connectivity layer, but there is an on-going debate about which network types and protocols are better suited for supporting mass sensor deployments. The decision regarding which network is to be used is an early consideration for IoT deployments, but how do you determine availability and which network will give you the best solution and project outcomes?

The short answer can be captured in two words, evolution and innovation. Over time, IoT sensor technology has evolved in capability, potential for scale and reduction in cost per module, creating a demand for new wireless network protocols and methods to support the new sensor types, many of which rely on battery power and infrequent messaging at long range, and over a wide area.

Doubt, uncertainty and fragmentation in the IoT market, combined with increasing sensor hardware and software innovation, have led to the creation and fielding of

several network connectivity options, each with their own attributes.

In recent years we've witnessed the growing expansion and evolution of one dominant IoT Network type, the Low Power Wide Area Network (LPWAN). The early entrants to this market, LoRaWAN and Sigfox, use free-to-air radio spectrum, and have had time to establish themselves across the world. LoRaWAN in particular has been a runaway success as an IoT Network connectivity option, dominating the market with over 40% market share of new connections, which is projected to continue adding market share through 2025.

On the cellular side, for new IoT network protocols NB-IoT and LTE-M (evolutions of the 4G spectrum that have now been adopted under the 5G standard), there is still an element of catch-up in progress. The GSMA was late in ratifying the standards for these IoT protocols and ultimately their deployment by Tier 1 carriers came sometime after the initial rollouts of the first LPWAN network



Author details:
Nick Sacke,
head of IoT and
products at
Comms365

connectivity protocols.

Despite initial predictions claiming that the cellular IoT Network variants would dominate the IoT connectivity market and squeeze LoRaWAN and Sigfox to the margins, there has been a lack of intensity in UK rollout of the cellular IoT network programmes (at the time of writing, LTE-M has been enabled in the Eastern half of the UK, and NB-IoT has 'holes' in its' coverage, particularly in the Eastern side of the UK). This means that IoT cellular LPWAN work has been largely limited to testing in the UK, while production rollouts are dominated by Private Council LoRaWAN installations and innovation programmes on public network variants of both LoRaWAN and Sigfox.

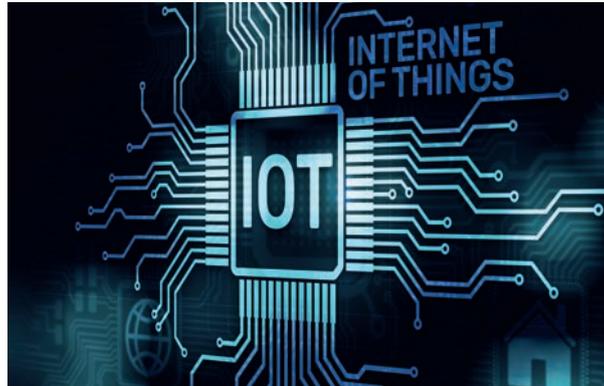
Globally, analysts project that there will be a 50:50 split in LPWAN network deployments between the free-to-air (unlicensed spectrum) and cellular variants (licensed spectrum) – the competition between these network standards will continue for some time to come. In particular, once 5G is fully rolled out and there are radio modules at a workable cost point for IoT, 5G protocols will also have an IoT element for the cellular side that will bring increased scale and efficiency in terms of its capability to connect millions of sensors per square kilometre.

Which network type or protocol to use?

As with most projects, cost of delivery for data is a primary concern that must be addressed. LoRaWAN and Sigfox are now at a level of maturity where the devices are cost-effective. Initially, cellular was a much higher cost, but is now starting to achieve cost-effectiveness. But in terms of usage costs, for NB-IoT and LTE-M, users are still paying for data usage on the network (paying by the byte), whereas LoRaWAN leverages the free-to-air spectrum facility and charges are based on device licensing, and in the case of Sigfox, per message.

Even though there appears to be a clear differential in terms of cost models, the choice of network and protocol isn't straightforward. As IoT rollouts become more commonplace, there are elements within a LoRaWAN environment that create cause for concern. With multiple devices sharing the LoRaWAN spectrum there can be potential collisions on the network and lost messages. In order to ensure each message arrives, the LoRaWAN protocol and software controlling the network has been adapted further to mitigate against this happening by spreading messages across multiple channels, monitoring message counters, and other techniques.

Identifying the parameters of the use case and the nature of the deployment is very important. If a message with telemetry data only needs to be sent when there is a status change, this won't necessarily create network congestion on a LoRaWAN Network. Regular 15 minute monitoring from multiple sensors in an area may however require additional gateway capacity to ensure spreading the sensor message load. But if your use case requires guaranteed delivery of traffic within a specific time period, or a constant stream of messages cellular protocols such as NB-IoT and LTE-M may need to be used.



Use cases

Another consideration is the protocols certain sectors are already using to gain traction. There's been a tremendous uptake and interest in LoRaWAN among local government that see it as a mechanism they can use to scale multiple use cases at once.

In the utility monitoring sector, NB-IoT appears to be the protocol gaining the advantage. No gateways are required as signal towers are the enablement point and it has deep penetration under the ground with good signal strength to reach its destination. But when it comes to monitoring elements deep within buildings, LoRaWAN can be more effective compared to what NB-IoT can do from the outside in. Refrigeration and temperature monitoring is one such example and LoRaWAN can provide an effective protocol in this instance, measuring the temperature deep inside the building, all the way down to the probe where intensive monitoring and data collection is critical.

For a use case such as measuring readings intermittently, on an alert basis or once an hour, you need to conserve battery power so the sensors last a long time and will likely be leading towards the unlicensed spectrum. Whereas within a healthcare monitoring scenario, such as in someone's home or in an ambulance on the move, readings will need to be sent through immediately, so will need to rely on the licensed spectrum, such as LTE-M.

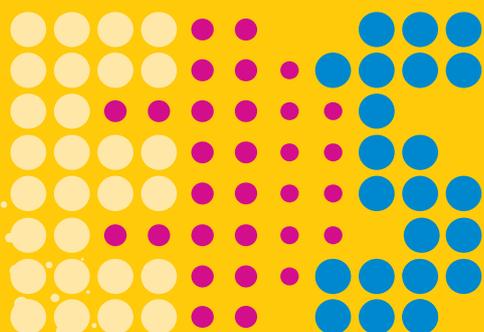
Blended connectivity

Currently, there is no one protocol that is optimised for every use case or can cover an entire estate. The solution is to deploy a hybrid model, one which blends different connectivity protocols together, from the unlicensed and licenced spectrums, to achieve total estate and use case coverage. A blended approach is inherently flexible, cost-effective and scalable for those looking to reap the benefits of mass-scale IoT but uncertain as to how and where to proceed.

Longevity is crucial for the success of an IoT deployment. No business wants to rip and replace the technology after ten years. It's clear that LoRaWAN in particular is on a growth trajectory that will provide that longevity. And the eventual maturity of 5G will also become another option for IoT projects, with much more efficiency in terms of capability to connect millions of sensors. 5G may be a way off yet, but it's likely that many estates of devices and networks will eventually have parts of them consumed by 5G.

With a blended model of different protocols covering each estate, to make it efficient and streamlined, it's important for a single platform to be used that can bring it all together and be received, read and analysed in one place. NB-IoT, LoRaWAN, LTE-M, Sigfox are all becoming industry-standard protocols that are each received in a different format. But they can be streamlined into one hub that intercepts the traffic and converts it into a protocol that the receiving application requires. By working with a partner offering all types of IoT connectivity in a blended solution, projects can be rolled out in the confidence that each protocol has been considered and is supported, to maximise the functionality, practicality and cost-efficiency of an entire IoT project.

31 MARCH _ 1 APRIL 2021



ENGINEERING DESIGN SHOW

> RECONNECT

THE NEW VIRTUAL EXPERIENCE FOR THE ENGINEERING DESIGN SECTOR

> A brand-new, two-day virtual event featuring a series of inspirational keynotes, case-study led presentations and interactive panel sessions, all addressing the key challenges and opportunities facing the UK engineering sector in 2021.

> EDS Reconnect will showcase the latest cutting edge technology, projects and innovations spanning the length and breadth of UK engineering.

> The best part... you can join us wherever you are and it's all free to attend.

REGISTER
NOW

[RECONNECT.ENGINEERINGDESIGNSHOWS.CO.UK](https://reconnect.engineeringdesignshows.co.uk)

 @EngDesignShow • #EDSReconnect  EngineeringDesignShow

Interested in exhibiting?

Head to the website to find out more about our virtual booths and sponsorship opportunities.

HEADLINE SPONSORS


Five Years Out


AHEAD OF WHAT'S POSSIBLE™


memory and storage specialists

Arrow in partnership with Analog Devices

EST. 2009

b'eeea

british engineering excellence awards

CELEBRATING DESIGN EXCELLENCE

JOIN US AS WE UNVEIL OUR 2020 WINNERS!

Thank you to all that entered this year's British Engineering Excellence Awards. Paul Fanning, publishing director at MA Business, said: "The British Engineering Excellence Awards have been the benchmark for great British engineering design for more than a decade now. I'm delighted to say that - despite all the disruption we've seen - this year is no exception, with an exceptionally diverse and fascinating set of entries".

To see who made the 2020 shortlist, visit: www.beeas.co.uk/finalists

The winners of this year's awards will be revealed during British Engineering Excellence Week, a special week-long programme of virtual festivities that will be hosted on the New Electronics and Eureka! websites from **22-26 March 2021**.

Keep an eye on the BEEAs website for more details and to find out how you can get involved!



 @TheBEEAs #BEEAS2020

WWW.BEEAS.CO.UK

HEADLINE SPONSOR

DISTRUPOLSM
A Univar company

CATEGORY SPONSORS

 COMSOL



SOLIDWORKS

Goodfellow
Your global supplier for materials

ORGANISED BY

newelectronics
Eureka!



MA BUSINESS
A MCB GROUP COMPANY