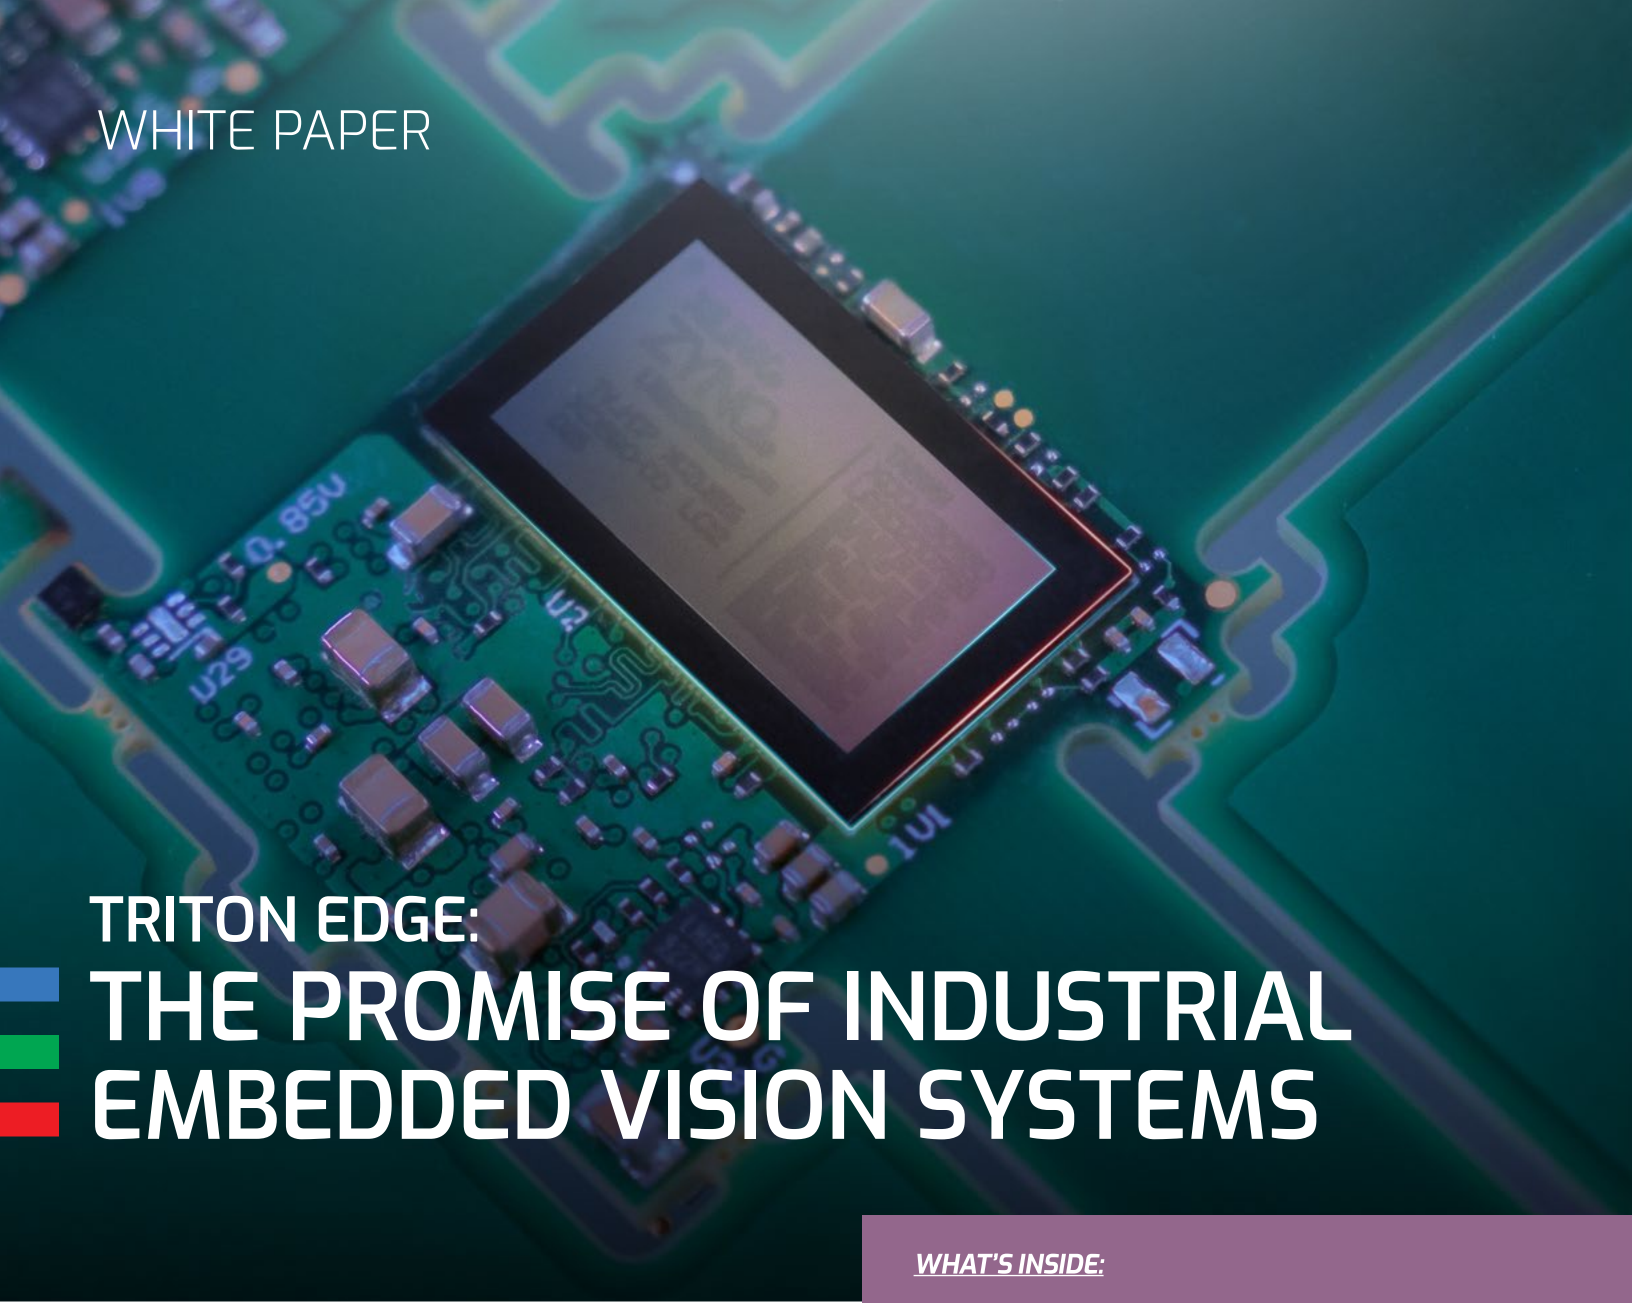

WHITE PAPER

TRITON EDGE:
THE PROMISE OF INDUSTRIAL EMBEDDED VISION SYSTEMS

This white paper will discuss how LUCID's **Triton Edge camera** helps vision application designers reduce their time-to-market while integrating their own IP into a compact vision system. By offering an innovative industrial IP67 camera powered by Xilinx'® Zynq® UltraScale+™ MPSoC (Multi-Processor System-on-Chip), LUCID effectively removes many of the steps needed to design and manufacture a compact embedded vision system. Validated to withstand the hardships of industrial use, the Triton Edge allows application designers more time to focus on creating their own innovative vision processing.

WHAT'S INSIDE:

LUCID VISION LABS

MOVING TOWARDS EMBEDDED

For decades, machine vision cameras have been a key component in driving factory and industrial automation. Since the beginning, machine vision cameras streamed images to host x86 PC computers (Intel/AMD 32/64 bit processors) for processing and analysis, allowing for an exponential jump in inspection speed and accuracy versus manual human inspection. While cameras and PC hardware both continued to improve in quality and processing power, this paradigm of a streaming camera connected to a host x86 PC remained mostly unchanged. The broad processing power and general purpose use of an x86 computer allowed application designers more flexibility in their vision processing, however at the expense of application size and power consumption.

In recent times, applications such as mobile or handheld systems created a need for smaller PCB sizes, lower power consumption, and more specialized processing. These challenges motivated factory automation application designers to seek chip makers with embedded platforms where specialized CPUs, peripheral devices, communication interfaces, and power management systems are all placed on a compact integrated circuit, (instead of separate components placed on a large motherboard in a typical PC). This single integrated circuit is known as a System-on-Chip (SoC) and ushered in an era of ultra compact computers. Many machine vision camera manufacturers have been quick to support this technology with compact, low-power cameras. However, to truly take advantage of embedded hardware technology and its promise of miniaturized, low power operation, application designers not only need to work through many time-consuming steps to reach a finished product but also need to create unique vision processing IP that customers value. This is a tough challenge but companies that can quickly and successfully reach their market before their competitors stand to benefit in larger growth and market dominance.

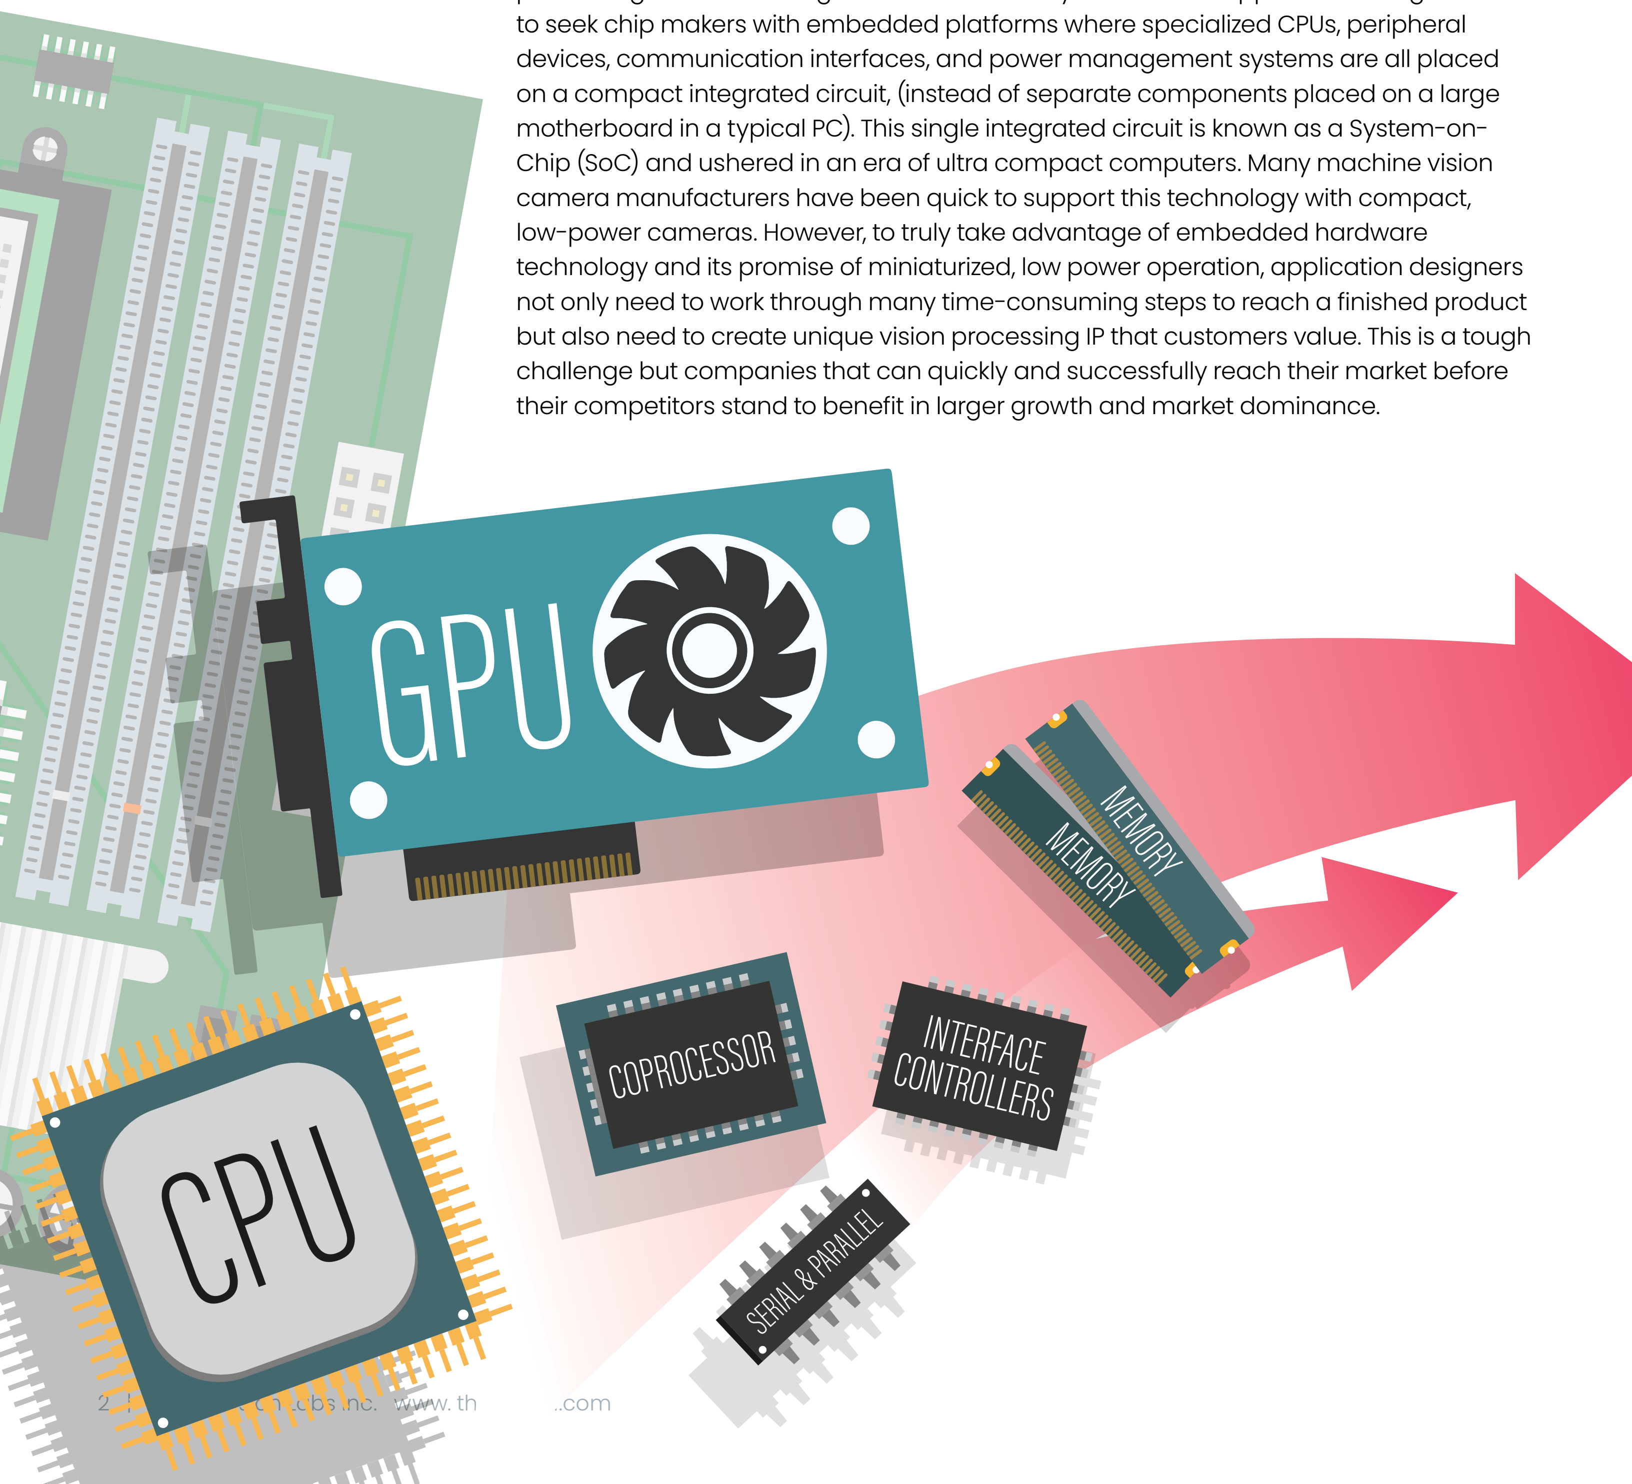

WHERE TO START?

1) EMBEDDED DEVELOPMENT KITS

While SoCs lack the general purpose processing power of an x86 PC, they allow for vision application designers to build smaller systems, performing only key specific tasks with greatly reduced power consumption and size. SoCs mostly use some version of a single or multi-core ARM processor ***(figure 1)*** for application processing. They also use a GPU for video processing and interface controllers such as USB, Ethernet and MIPI are also included in the SoC. Shared system RAM, power management systems, eMMC storage and any other unique processing units (i.e., an AI DPU) are soldered on a larger PCB called a system-on-module (SoM) alongside the SoC. The SoM then connects to a "carrier board" which held the physical expansion and interface ports for peripheral connections ***(figure 2)***.

To promote their SoC and SoM technologies, major chip manufacturers like NVIDIA® offer development kits or evaluation kits which package their SoCs or SoMs with a general-purpose carrier board equipped with a variety of interface ports. In some cases, the SoC or SoM would be directly soldered into the carrier board, making it a Single-Board Computer (SBC). These development kits allow vision designers a quick way to evaluate the chip maker's embedded technologies. Current examples of embedded development kits that are popular with vision system designers include kits from NVIDIA (Jetson Nano, Jetson TX2, Jetson Xavier NX), Xilinx (Zynq 7000 SoC, Zynq UltraScale+ MPSoC), and the Raspberry Pi Foundation (Raspberry Pi 4 (SBC)). In addition to these boards, multiple 3rd party hardware vendors offer unique carrier boards that provide a range of different hardware features such as additional peripheral ports and other hardware expansion slots. For example, some carrier boards may provide additional MIPI connectors, GPIO ports, PCIe, or more memory slots compared to the original development kit.

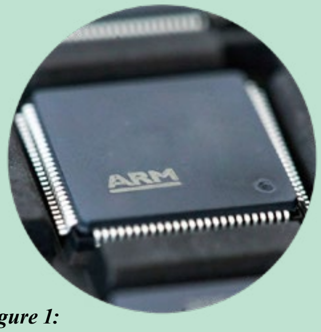

figure 1:
ARM Processor Chip

ARM processors are based on a reduced instruction set computer architecture (RISC). This architecture enables the processors to operate at very efficient performance per watt levels. In general, ARM processors can be extremely fast but require more effort on the programmer's part to design high-quality instructions.

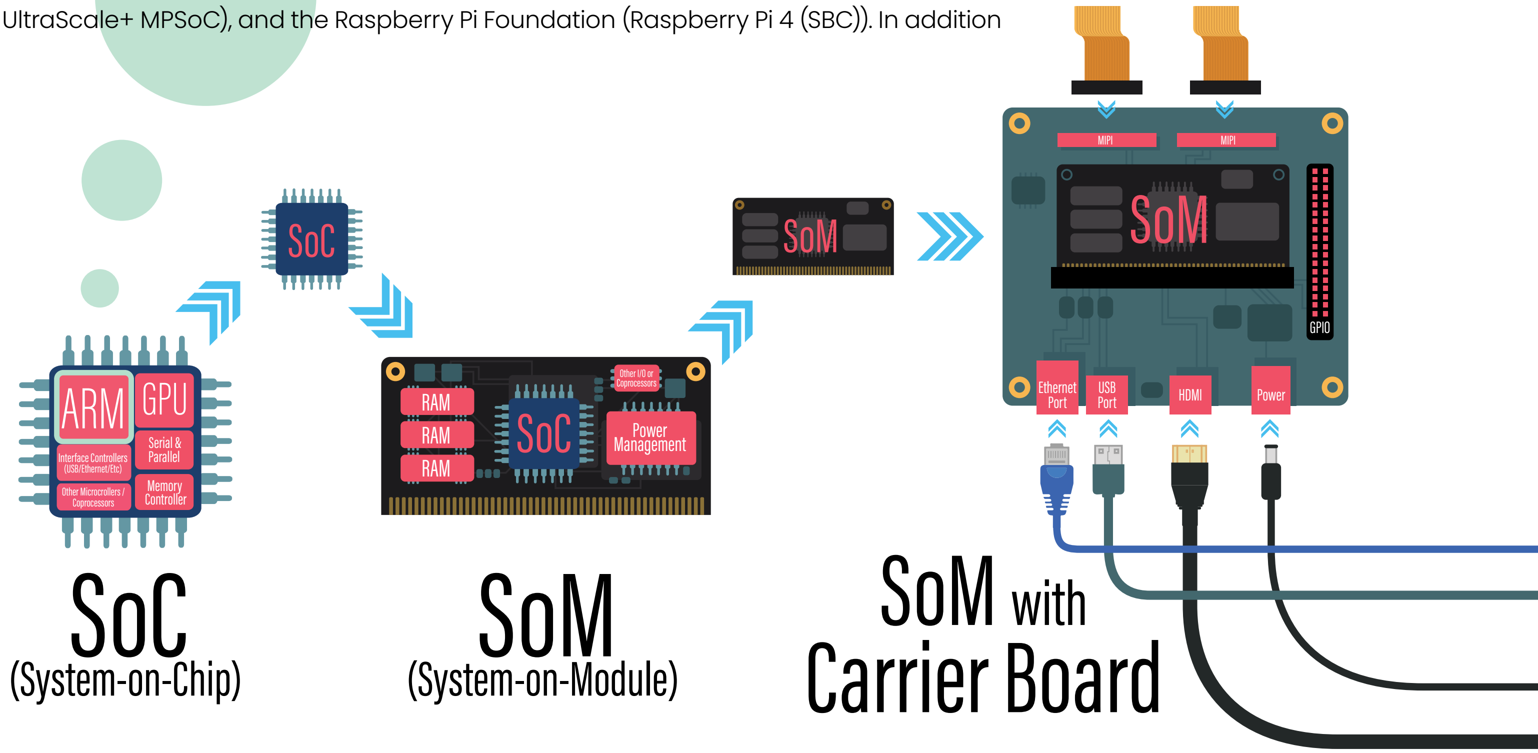

figure 2: *SoC, SoM, and Carrier Board*

The basic building blocks of an embedded development kit. A specialized chip known as an SoC is placed on a specialized module, known as a SoM. This SoM then connects to a carrier board which contains physical ports for peripheral connections.

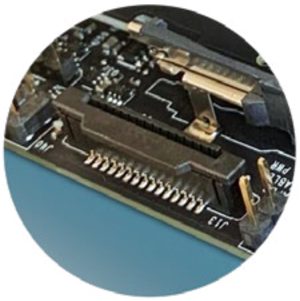

figure 3: *MIPI Connection*

For the purpose of connecting peripherals such as external displays and cameras, the MIPI interface standard became the preferred interface. This interface allows for chip-to-chip communication and includes a low-power camera serial interface (CSI) for camera modules.

2) EMBEDDED VISION CAMERA MODULES

While traditional machine vision cameras can be connected to a USB or Ethernet port on the carrier board, those ports are typically limited in number, and connecting a GigE camera to the Ethernet port effectively removes that interface for other connectivity uses. Instead of using a traditional machine vision camera, camera modules connect via the MIPI port. These camera modules usually feature low-power, low-cost, and board-level (no case) rolling shutter sensors and provide a compact alternative to larger machine vision cameras ***(figure 3)***. However, for most of these camera modules a camera adapter board or camera expansion board is needed. The adapter board or expansion board is a camera manufacturer specific PCB that sits between the camera module and the development board's MIPI connection or camera expansion port. These additional boards bridge any differences between the pin layout and the data lanes between the camera and the development board. In some cases, the adapter or expansion board will feature a separate power supply for the camera. While the camera module itself is much smaller than a traditional machine vision camera, the added need for an adapter board, expansion board, and additional MIPI and power cables quickly increases the overall size.

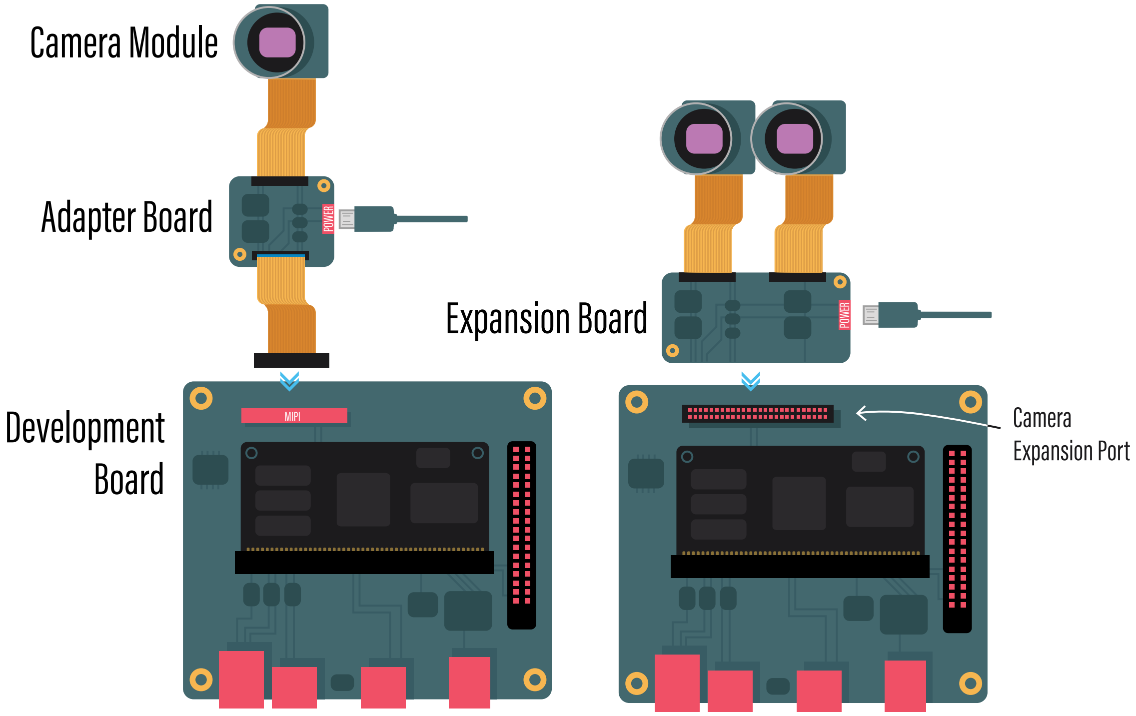

Early camera modules also came with very little on-board image processing, opting to push most, if not all the basic image processing onto the SoC's processors. This image processing would take up valuable resources on the embedded board and would limit the amount of additional vision processing that engineers could create on the system. Soon after however, machine vision camera manufactures would take their camera design expertise and build camera modules with an on-board FPGA or ASIC chip. These newer camera modules would be capable of running standard image processing such as auto white balance, auto exposure, auto gain, region of interest (ROI), debayering, and more. While engineers couldn't modify or add additional processing on the camera, they could at least access and run these pre-installed processing features on the camera module itself, freeing up embedded board resources for other processing tasks.

MAJOR CHALLENGES

BALANCING EDGE PROCESSING TODAY

Current embedded vision application designers are creating more customer value by utilizing both the camera's on-board image processing and by adding unique vision processing to run on the embedded development board. Compared to a modern Intel/AMD PC, development boards are more limited in resources and application designers need to strategically balance available resources between components. If a balance can be found and a proof-of-concept (POC) is viable, the POC can be prototyped into more portable applications such as kiosks, aerial drones, and autonomous robotics where it is beneficial for edge processing (localized processing that happens on the embedded vision system instead of the host PC, server, or cloud.) This is also true in the industrial space, where vision systems using embedded hardware technologies are benefiting from streamlined designs that offer custom image processing such as AI inference for object detection and classification, along with a smaller physical footprint and reduced power consumption.

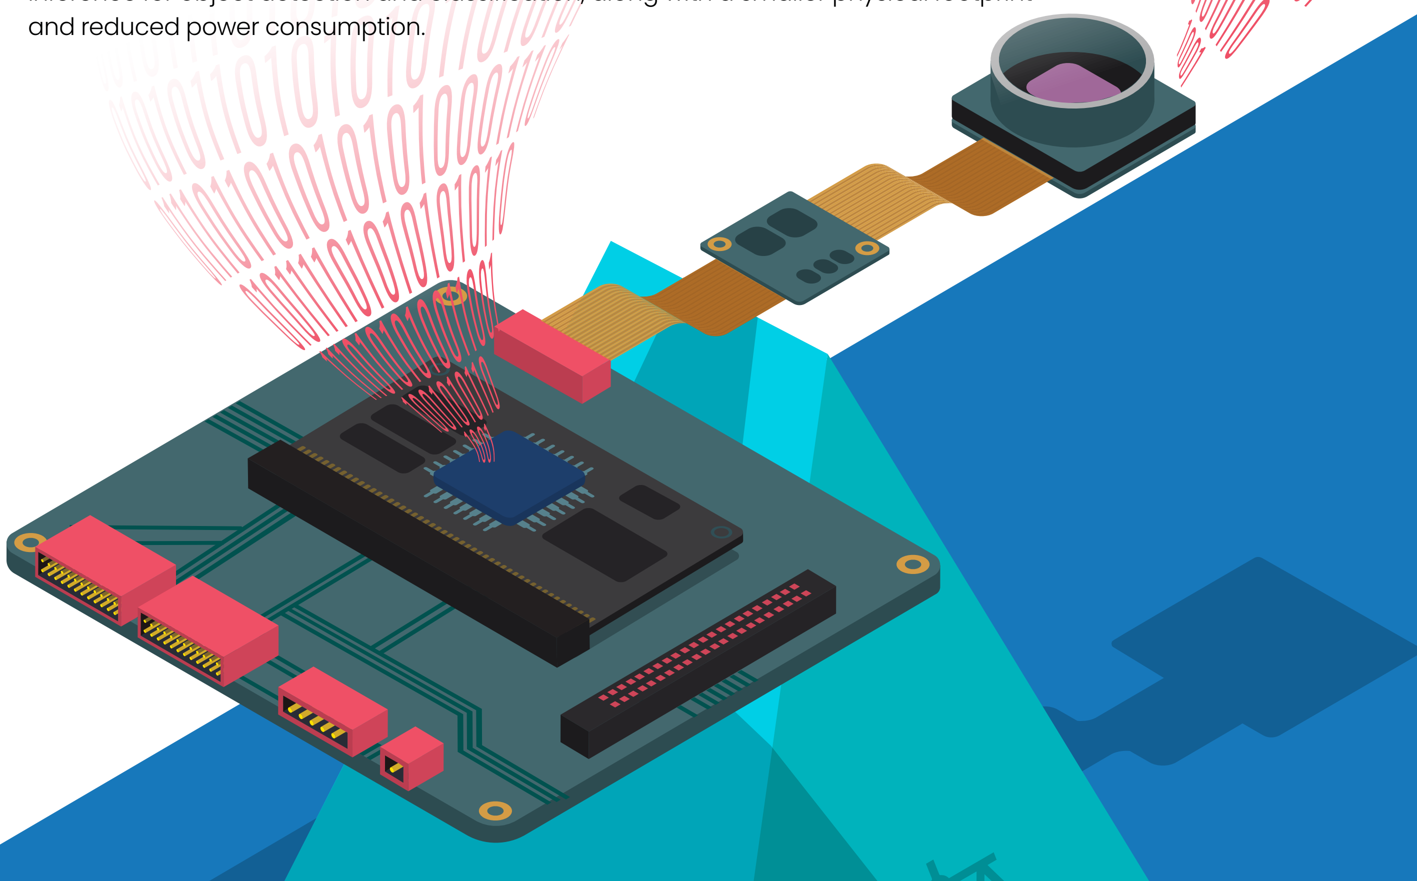

SURVIVING INDUSTRIAL CHALLENGES

To build these compact embedded vision systems however, application designers must navigate the challenges of harsher operating environments and the complexities of building smaller, faster, more power-efficient systems. They must work to validate their system through time-consuming stages, starting from the proof of concept, to prototyping, and finally to a minimum viable product (MVP) or a Full Custom Design (FCD). Off-the-shelf embedded development kits, such as those from NVIDIA, Xilinx, or Raspberry Pi, offer a quick solution to building a proof-of-concept design. However, many camera modules and embedded development boards offer little to no protection from the harsh environments of industrial spaces. A considerable amount of time must be spent on designing and testing prototypes that are protected against dust and moisture (IP67 or IP65), electromagnetic interference (EMC immunity) and physical shocks and vibration (random and sinusoidal) – common stressors not only in industrial spaces but also in commercial and portable mobile spaces.

Moreover, developing a lean and miniaturized end-product with custom silicon from a proof-of-concept demands a high level of expertise and time. Once a proof-of-concept is validated using a development board, vision engineers must strip away unnecessary components to design and build an optimized application that takes up minimal space. This process requires a custom carrier board that utilizes the original SoC or SoM with only the very necessary components of the development board. Unless the vision application designers have their own SMT lines for custom PCB manufacturing, this stage would require a 3rd party hardware manufacturer.

While embedded development boards can offer time savings for application designers during the PoC stage, and in some cases the prototyping stage, there is still a considerable amount of time, resources, and costs needed to reach a miniaturized MVP or FCD that can function in an industrial setting. In today's competitive vision system environment, development stages that take too long risk losing ground to more agile rivals.

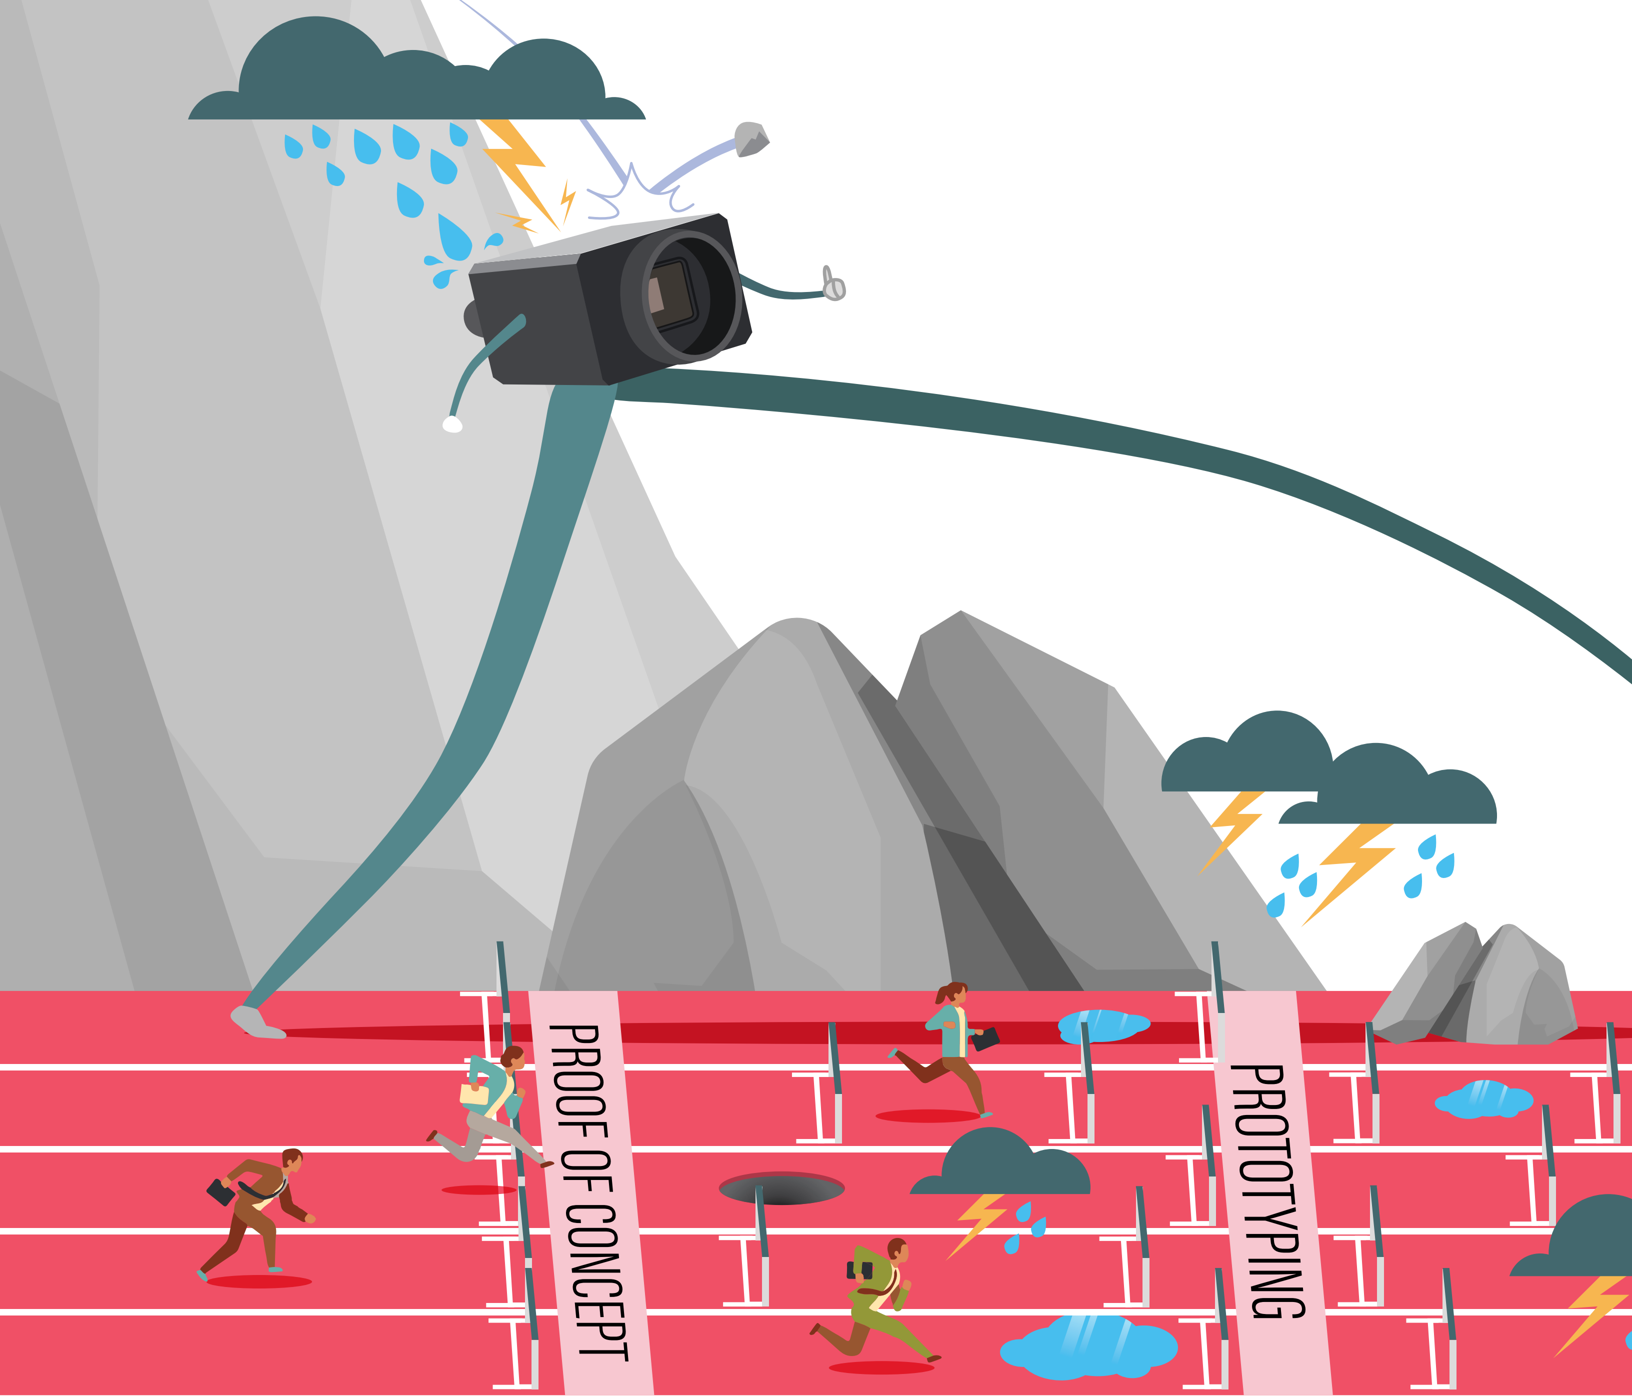

figure 4: *Triton Edge industrial certification can help engineers skip time-consuming and costly hardware validation steps. The camera is certified with:*

Shock and Vibration
• EN 60068-2-27
• EN 60068-2-64
• EN 60068-2-6

Industrial EMC Immunity
• EN 61000-6-2

Dust Proof, Water Resistant
• IP67 Rated

M12 and M8 connectors
• IEC 61076-2-109
• IEC 61076-2-104

ALL-IN-ONE EDGE COMPUTING CAMERA

LEAPFROG OVER HARDWARE DEVELOPMENT TIME

LUCID's solution to simplifying these hardware development stages is by integrating Xilinx's powerful embedded chipset, the Zynq UltraScale+ MPSoC, into the Triton Edge camera itself. To realize a faster time-to-market, vision application designers can skip many of the hardware validation steps needed to qualify their product for challenging environments thanks to the Triton Edge's "Factory Tough" design.

The compact and lightweight aluminum camera provides IP67 protection, is certified against physical shocks and vibration, offers EMC industrial immunity, and provides an operating temperature between -20° C to +55° C ambient ***(figure 4)***. LUCID and Xilinx have also done the work to miniaturize the entire camera. Thanks to Xilinx's compact inFO package chipset size ***(figure 5)*** and LUCID's flex-ridged PCB camera design, the Zynq UltraScale+ MPSoC's twin dual-core ARM processor units, user-programmable FPGA (Field Programmable Gate Array), RAM, and storage all fit in an ultra-compact camera measuring only 29 mm x 44 mm x 45 mm (H x W x D). The Triton Edge allows for high-speed video direct memory access (AXI VDMA) between the on-camera Image Signal Processor (ISP), user-programmable FPGA, and on-board RAM. The ARM cores also utilize their own DMA engine, freeing up the processors from managing data transfers. The VDMA and DMA reduce system bottlenecks, frame buffer overhead, and memory access latency allowing for vision processing to run as efficiently and as fast as possible. All of this translates into a hardware optimized, pre-miniaturized industrial embedded camera system ready to go out-of-the-box.

figure 5: *Zynq UltraScale+ MPSoC in InFO package size*

InFO devices are 60% smaller, 70% thinner, with better thermal dissipation and higher signal integrity, all without sacrificing the processing power of the Zynq UltraScale+ MPSoC.

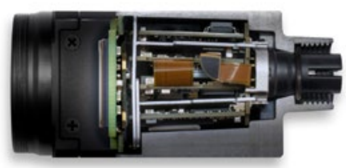

EDGE PROCESSING POWER FOR YOUR VISION IP

With this compact industrial camera in place, application designers will have more time to develop on-the-edge processing running on the camera. Unlike traditional machine vision cameras that need to transfer image data to a host PC before any specialized processing can begin, edge processing can take place inside the camera itself, without the need to connect to an embedded board or any type of host PC. Edge processing allows an embedded vision system to process time sensitive image data with less latency compared to host PC, server, or cloud-based image processing. It also reduces network usage, allowing the system to only send the processed results. For example, instead of the camera sending multiple images with different exposures to a host PC to create a high dynamic range (HDR) image, it can send just the final processed HDR image.

As already mentioned, creating unique vision processing IP that customers value is a major challenge. Having access to a wide range of edge processing resources is necessary to realize the full potential of a designer's goals ***(figure 6)***. With the Triton Edge camera, LUCID provides a pre-installed ISP that offers many of the same features of a regular machine vision camera's ISP (auto white balance, auto exposure, auto gain, gamma, CCM, etc). The ARM Cortex-A53 processor can run application tasks such as Linux-based programs. The ARM Cortex-R5 can be utilized for real-time tasks such as motor control. A user programmable section on the FPGA is available for engineers to create more deterministic and unique vision IP, allowing them to reach faster performance levels compared to traditional software-based solutions. And finally, it offers 2GB of DDR4 RAM and 8GB of eMMC storage. This ability to access the FPGA is very powerful and when combined with multi-core ARM processors, LUCID's ISP, RAM, and storage, vision engineers have more flexibility than ever to create their own IP in such a compact industrial camera.

figure 6: *Zynq UltraScale+ MPSoC in the Triton Edge includes the following:*

Chipset	Xilinx UltraScale+ ZU3CG MPSoC
Processors	Dual-core ARM® Cortex™-A53 APU Dual-core ARM Cortex-R5F
System Memory	2 GB DDR4
Storage	8 GB eMMC

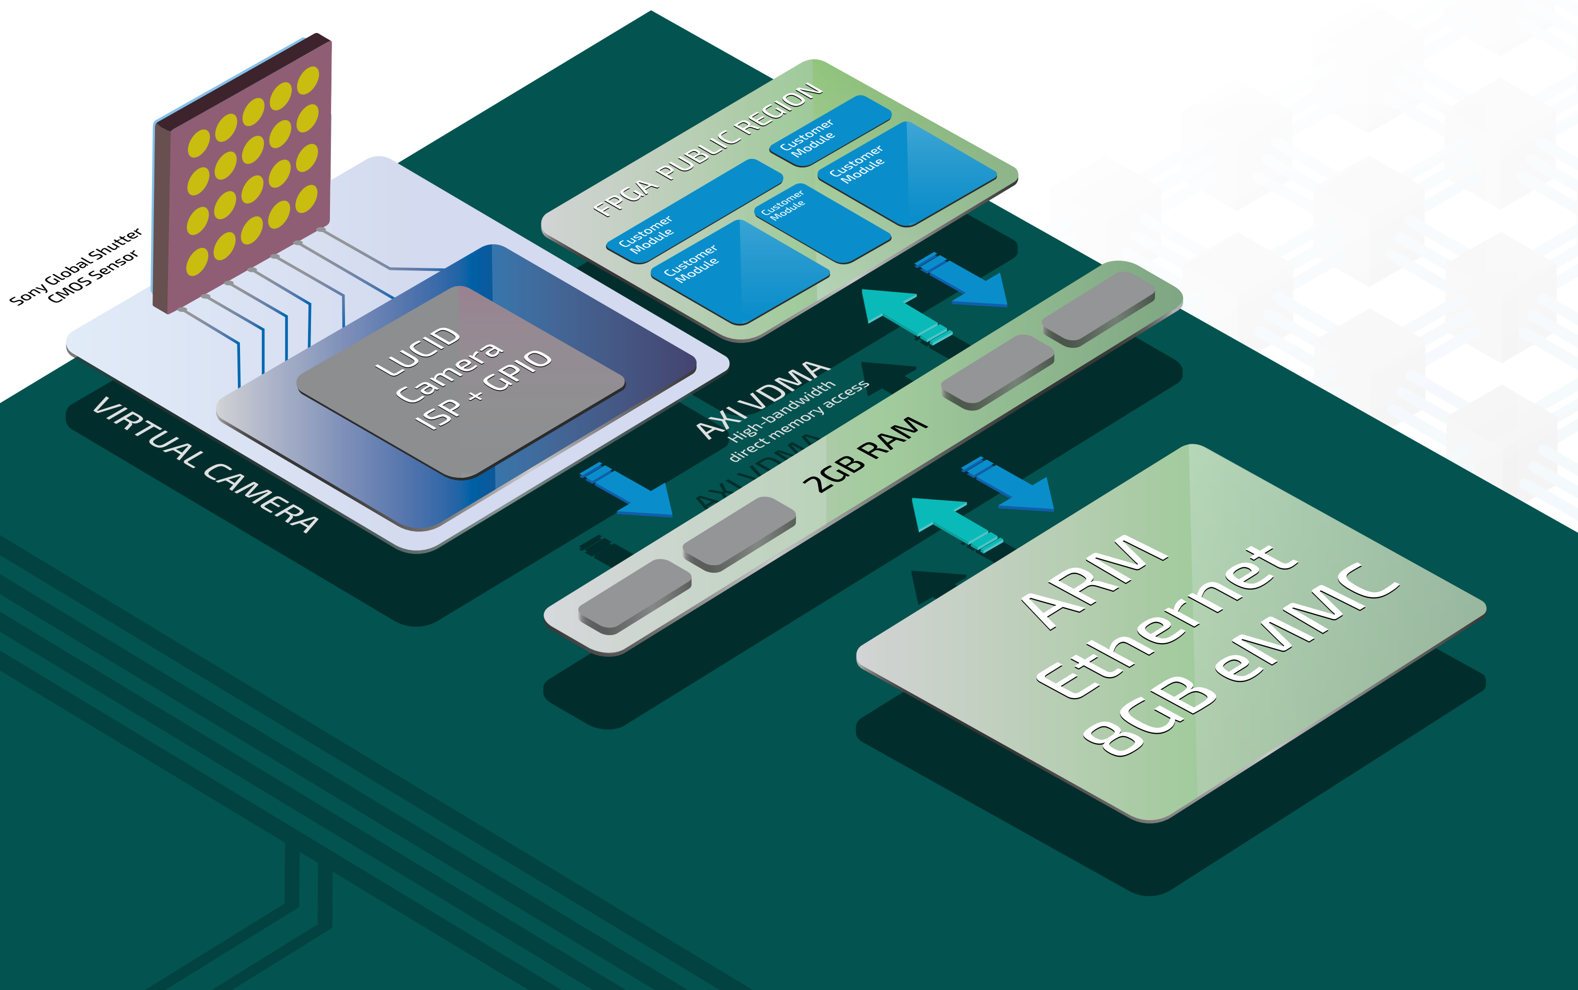

THE FPGA: SPEED vs POWER CONSUMPTION vs FLEXIBILITY

To achieve a compact size system, application designers must weigh processing speed against power consumption and future upgradability. The more power needed, the larger the device needs to be to dissipate heat. For example, integrating GPUs for accelerated parallel processing can be extremely powerful yet very power hungry, needing a larger overall size with heatsinks or active cooling fans. In addition, being able to upgrade hardware-accelerated vision processing on a finalized hardware design can be very challenging. Custom ASIC processors, for example, cannot be modified without doing a complete redesign of the silicon and for most vision applications designers, ASIC technology is out of the question due to their extraordinary high costs. The other option is to use the ARM cores for processing. However, any complex processing needing large multiple instructions to be executed at once (parallel processing) could be interrupted by other background processes (i.e., from the OS) running on the ARM cores. For a powerful balance between processing speed, power consumption, and future modification flexibility, the Triton Edge offers access to a user programmable region on the FPGA.

FPGAs provide more deterministic hardware-accelerated processing compared to ARM processing, lower power usage than GPUs, and without the high cost and development time of designing a custom ASIC processor. FPGAs consist of many hardware blocks with reconfigurable interconnects that can be changed at any time to implement almost any type of functionality ***(figure 7)***. FGPAs can handle multiple, specific instructions on each clock cycle, allowing for complex software algorithms to run exponentially faster with very low power consumption. Any future algorithm improvements or new processing features can be added to the camera's FPGA as well, helping one to further differentiate their product and services after deployment in the field. This level of hardware-accelerated flexibility allows for a more agile development environment where embedded vision designers can evolve processing to meet the changing needs of their users down the road.

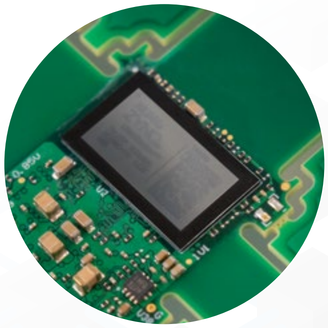

FPGA benefits:

- Provide deterministic, low latency real-time processing.
- Low power consumption.
- Operations can run in parallel.
- Long life cycle thanks to it being re-configurable.
- Cost-effective technology for more agile development.

figure 7: Graphical representation of FPGA

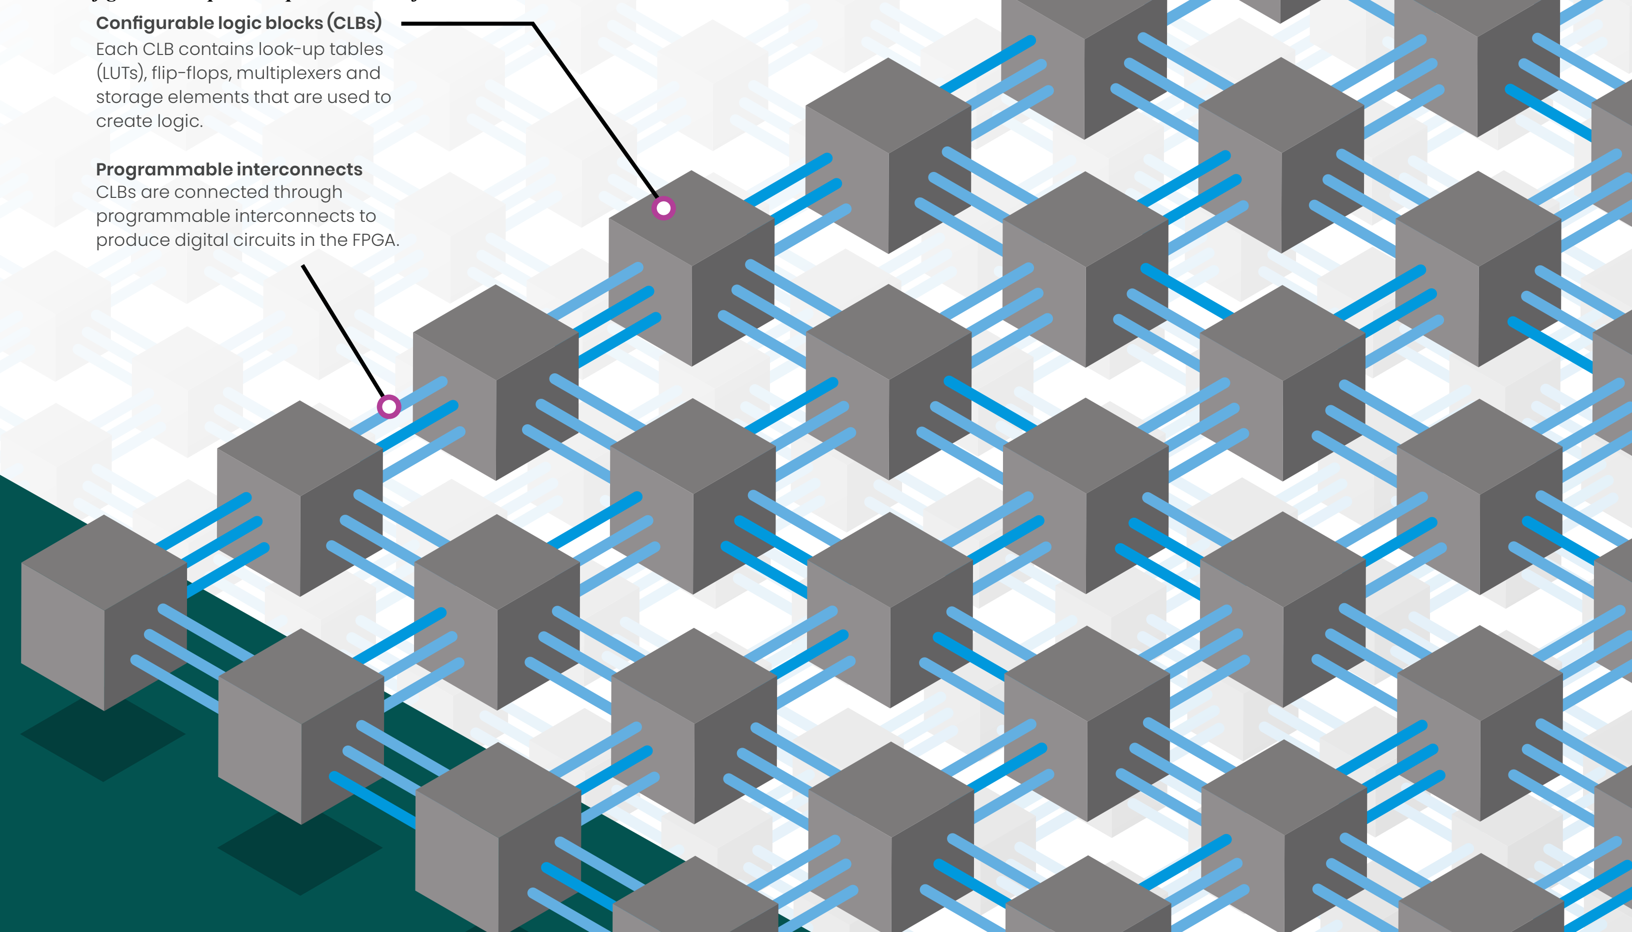

SCALING UP YOUR VISION IP & MULTI-CAMERA FLEXIBILITY

The system resources found in each Triton Edge allows one to scale-up the number of image streams in an application without having to worry about diminishing processing resources. Unlike MIPI camera modules, where adding additional modules increases the strain on the shared SoC and SoM resources, each Triton Edge is equipped with its own Xilinx Zynq UltraScale+ MPSoC. With the Triton Edge there are no resources needing to be shared or allocated between cameras. Scaling up by adding more cameras in a multi-camera application is easier, where each camera can do its own unique vision processing without taking away resources from others.

In addition, there is almost no limit to the number of Triton Edge cameras one can use, the only real limit is the number of Ethernet PoE ports on a network switch. Compared to MIPI camera modules, there are no physical carrier board or camera expansion board limitations, such as the number of MIPI ports available. Also, adding MIPI camera modules to an already finalized design would require a redesign for the carrier board or camera expansion board. Because the Triton Edge is already validated against industrial use, it can reduce future development time for when the need arises for increasing image streams in a multi-camera application. In the competitive space of machine vision, the Triton Edge allows engineers to quickly react to other challengers. Thanks to the camera's independent processing resources, engineers can upgrade their systems without struggling with resource allocation.

Multi-camera friendly
Add many Triton Edge cameras to an application without having to worry about system resource allocation or redesigning a custom carrier board.

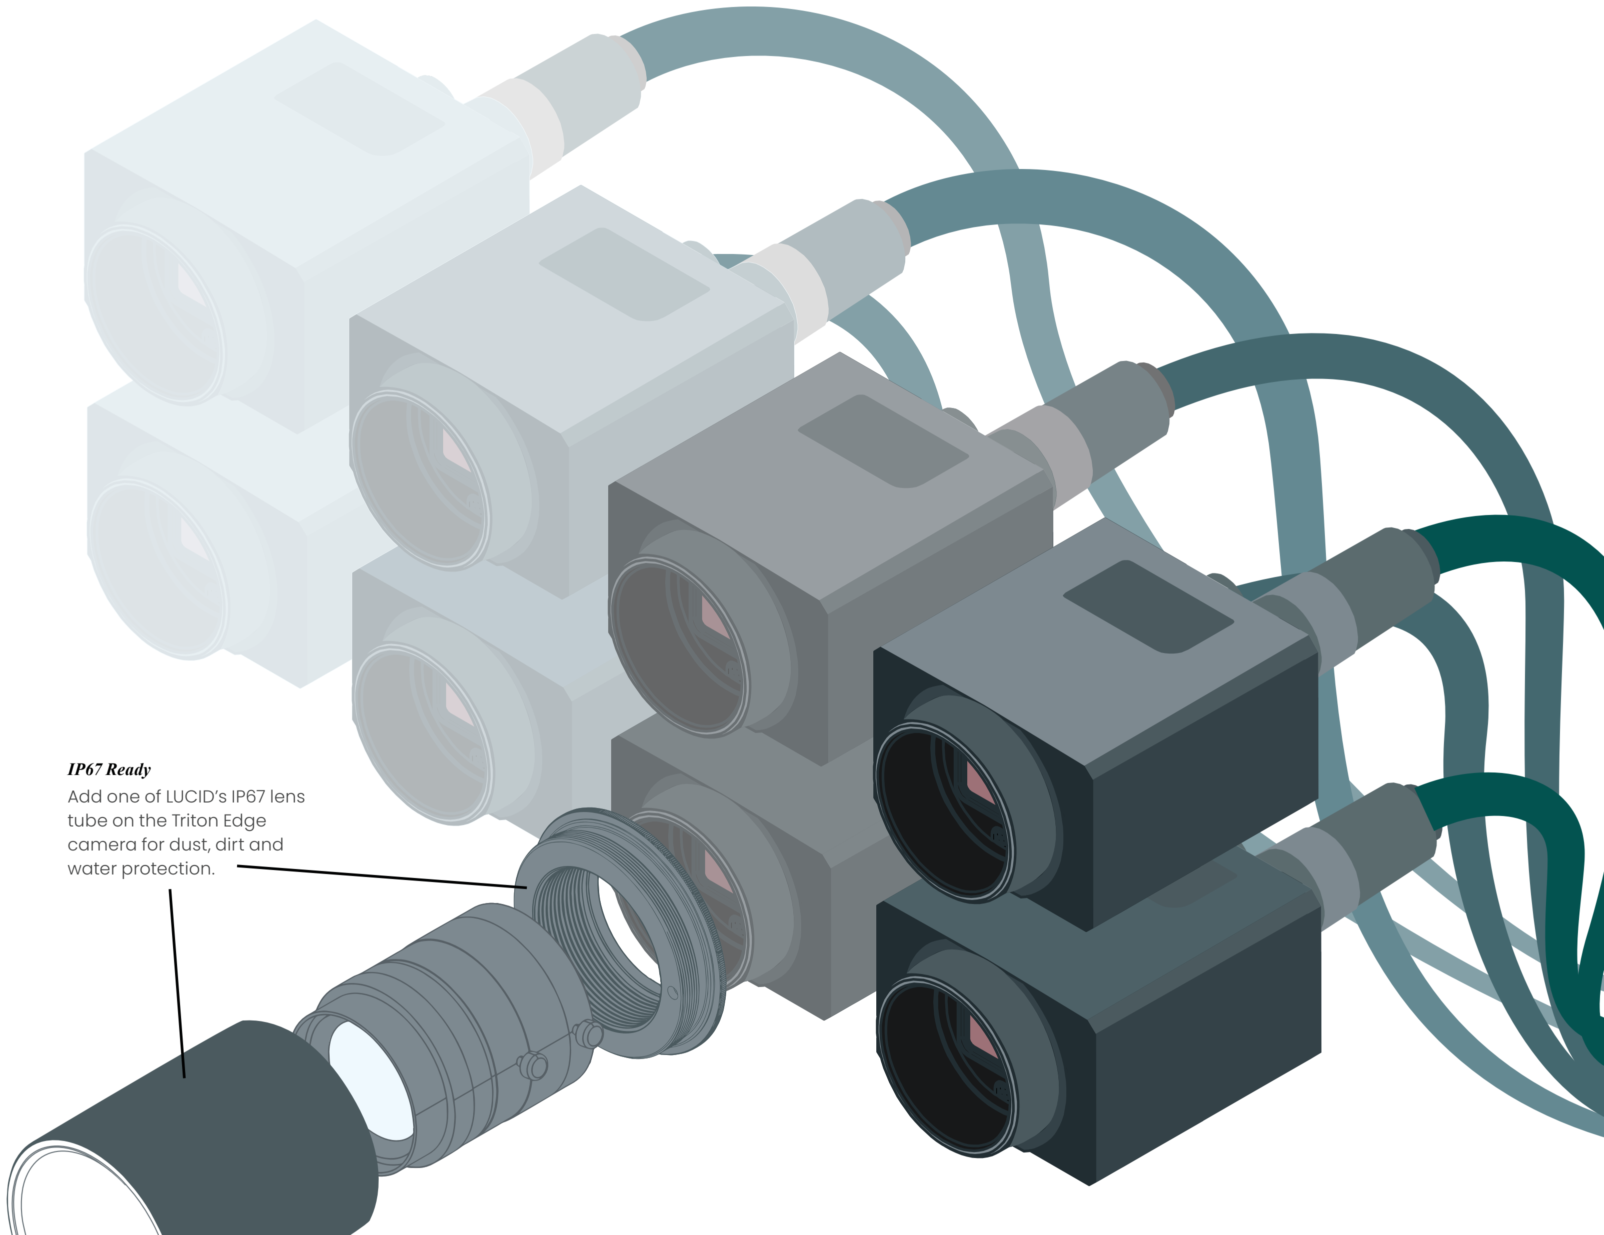

IP67 Ready
Add one of LUCID's IP67 lens tube on the Triton Edge camera for dust, dirt and water protection.

TOOLS TO BUILD YOUR VISION

Traditionally, machine vision camera manufacturers provide a Software Development Kit (SDK) that includes APIs to access the camera's feature set. These SDKs would typically include C/C++ and .NET libraries along with a graphical user interface (GUI) viewer designed to run on a Windows or Linux x86 host PC. While OEMs had the ability to create customized camera control software, they had to rely on the SDK's APIs running on the host PC. In addition, OEMs are locked out of the FPGA and had no way of creating FPGA-accelerated functions on the camera. With the Zynq UltraScale+MPSoC however, the Triton Edge provides almost complete control of the camera. Engineers will use the following combination of powerful software tools to build vision processing:

Xilinx's Vitis™ High-Level Synthesis (HLS) platform. Custom FPGA programmable logic, called "overlays" are created using Vitis. The unique Vitis development environment allows for C/C++ specifications to target directly into the user programmable FPGA region, without having to use ASIC-style design tools (Verilog or VHDL are still completely suitable if one chooses to use them). A number of reference designs are already available, including Xilinx's AI DPU. Store more overlays on the eMMC and swap them in and out of the FPGA for even more specialized processing.

PYNQ™ - "**Py**thon Productivity for Zy**nq** UltraScale+ MPSoC" – Open-source framework. PYNQ supports all major Python libraries like Numpy, Scikit-Learn, Pandas and others. Engineers can wrap custom overlays in PYNQ's Python API and end-users can access them via Python through the open-source Jupyter Notebook environment provided within the PYNQ framework.

Jupyter Notebooks. Jupyter Notebook provides a browser-based interactive environment for developing on-camera apps that can utilize custom overlays on the FPGA ***(figure 8)***. Create a custom interactive GUI where users can not only control the camera but also edit and run Python code in real-time without having to go through the traditional software compile stage. Jupyter Notebooks support live code, interactive widgets, plots, narrative text, equations, images, and video with complete access to PetaLinux file system on the camera.

PetaLinux. The Triton Edge runs a fully functional Linux Ubuntu OS giving users a familiar development environment. Open a Terminal window to download and install packages, such as your favorite compiler and build applications or APIs directly on the camera. For example, install OpenCV for ARM and create on-camera computer vision applications.

This combination of development tools can lead to some truly unique multi-stage processing. For example, additional custom overlays can be stored on the PetaLinux file system residing on the Triton's 8GB eMMC storage. Using Jupyter Notebook, these overlays can be switched in and out of the FPGA's public region at any time during image processing. Images simply stay in the 2GB of DDR memory while different overlay sets are loaded and applied. By utilizing these available tools, the Triton Edge brings powerful processing capabilities that allows to build a more bespoke vision solution for the end-user.

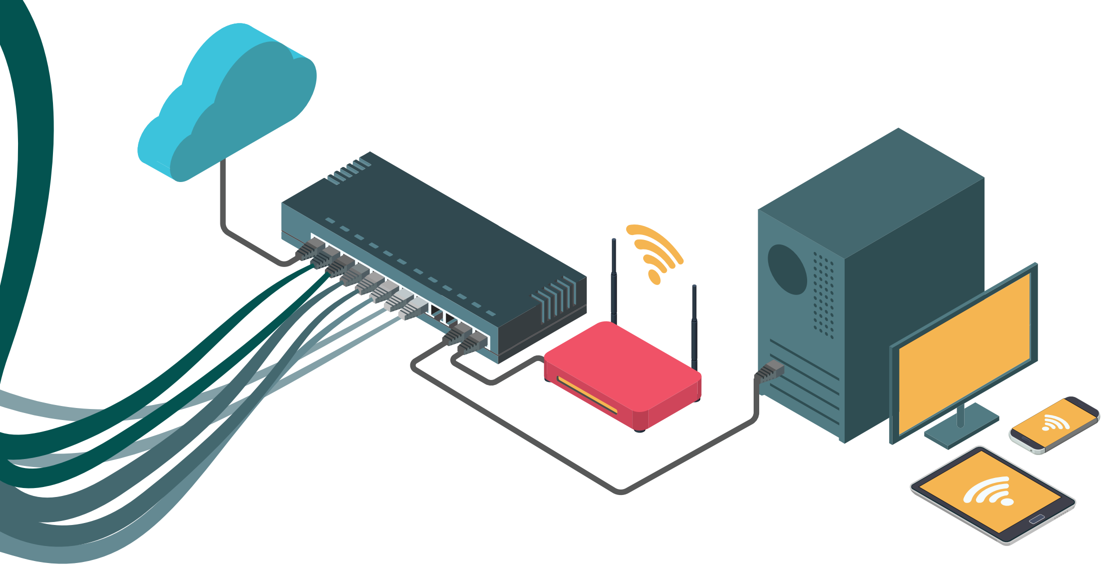

figure 8. Developing on the Triton Edge with Jupyter Notebooks.

Jupyter Notebooks running on the Triton Edge camera can be accessed through any device that can run a modern web browser.

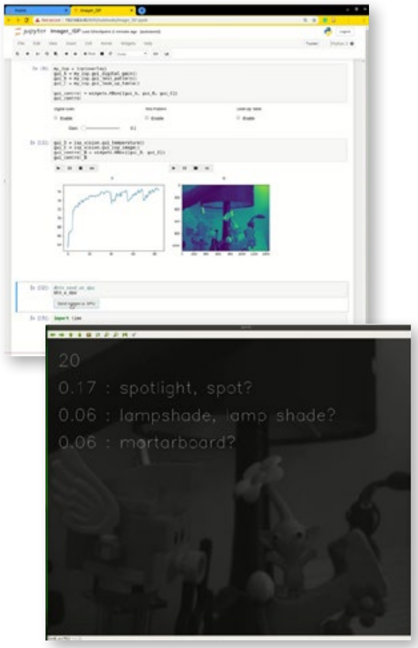

Example of Triton Edge running a Jupytor Notebook accessing the Xilinx Deep Learning Unit (DPU) running on FPGA during a live stream.

AI AND "EDGE-TO-CLOUD"

One of the more unique benefits to the Xilinx' Zynq UltraScale+MPSoC is its support for Vitis AI, Xilinx's AI inference development platform and Xilinx's Deep Learning Processing Unit (DPU) for Edge. Like other custom overlays, the DPU is built as an overlay and can be loaded onto either the FPGA's public region or stored on the eMMC for use at later times. The DPU is designed to work with many convolutional neural networks including VGG, ResNet, GoogLeNet, YOLO, SSD, MobileNet, FPN, and more. Users can deploy the Vitis AI development environment to accelerate AI inference on the Triton Edge, such as being able to quickly deploy models from TensorFlow and convert them to run directly on the DPU. In addition, the rapid move to distributed cloud processing can also be implemented on the Triton Edge. The Zynq UltraScale+ MPSoC allows one to build an edge-to-cloud device thanks to Xilinx's support for many of the popular cloud platforms. Support includes AWS IoT Greengrass and AWS SageMaker along with coming support for Microsoft Azure Sphere and Azure IoT cloud platforms. Whether one needs an AI DPU or cloud-connected camera, the flexibility of the Triton Edge and Xilinx's comprehensive support will help one rapidly react to changing market requirements.

Amazon SageMaker

AWS Greengrass

AWS Greengrass compatibility for Zynq UltraScale+ MPSoCs is built on the Xilinx PYNQ framework. AI models can be trained using Amazon SageMaker and deployed from cloud to camera using AWS Greengrass. Thanks to this, the Triton Edge can by designed as an edge-to-cloud device, collaboratively processing data between the camera and the cloud over intermittent connectivity.

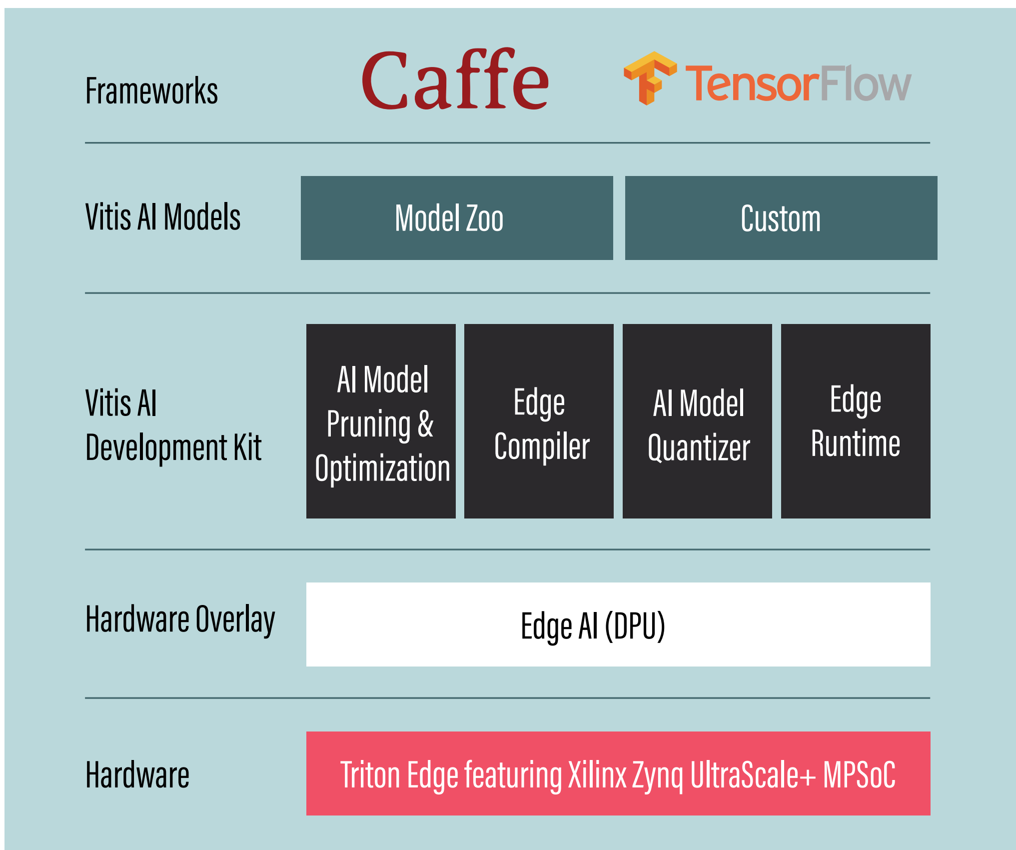

CONCLUSION: JUMP START YOUR VISION

LUCID's Triton Edge camera featuring Xilinx's Zynq UltraScale+™ MPSoC provides a new level of on-camera performance and flexibility without sacrificing power efficiency, sensor performance, or camera size. The camera's design allows application engineers to leapfrog the competition by providing faster time-to-market thanks to its miniaturized hardware design and validated industrial reliability. Powerful on-camera system resources, such as the twin dual-core ARM processors and FPGA, provide increased flexibility to develop unique vision IP while also allowing engineers increased scalability thanks to its carrier-board free design. Without having to rely on a camera manufacturer's SDK or the need to develop and run code on a host PC, OEMs are free to develop on-camera edge processing, including AI, through various development environments such as Xilinx Vitis, PetaLinux, and Jupyter Notebook. With the Triton Edge, application engineers can jump start their vision by cutting manufacturing costs and saving development time while providing more value to their end-user.

LUCID Headquarters
LUCID Vision Labs, Inc.
130-13200 Delf Place, Richmond B.C.
Canada, V6V 2A2
sales@thinklucid.com
thinklucid.com
1-833-465-8243

Europe, Middle East, Africa
LUCID Vision Labs GmbH
Renntalstraße 14, 74360 Ilsfeld
Germany
sales.emea@thinklucid.com
+49 (0) 7062 97676 12

Asia Pacific
LUCID Vision Labs G.K
Eishin Bldg. 4F 3-6-1, Kanda-Ogawamachi,
Chiyoda-ku, Tokyo 101-0052
Japan
sales.apac@thinklucid.com
+81 3 5577 7915

Greater China
LUCID Vision Labs, China
sales.gc@thinklucid.com
thinklucid.cn

LUCID VISION LABS