# Xilinx Expands into Wide Range of Vision Guided Machine Learning Applications with reVISION™

**BACKGROUNDER**
**March 13, 2017**

## MACHINE LEARNING: FROM EDGE TO CLOUD

Machine learning applications are rapidly expanding across a growing number of end markets, resident at the edge, in the cloud, or a hybrid where edge processing is combined with access to cloud-based data analytics. For the cloud, Xilinx recently introduced the Reconfiguration Acceleration Stack (launched November 2016), targeting a wide range of compute acceleration applications that include machine learning inference. For the edge, Xilinx is now announcing a significant expansion into a wide range of vision guided machine learning applications with the Xilinx® reVISION™ stack. The new reVISION stack enables a much broader set of embedded software and systems engineers, with little or no hardware design expertise to develop vision guided machine learning applications easier and faster with Xilinx technology.

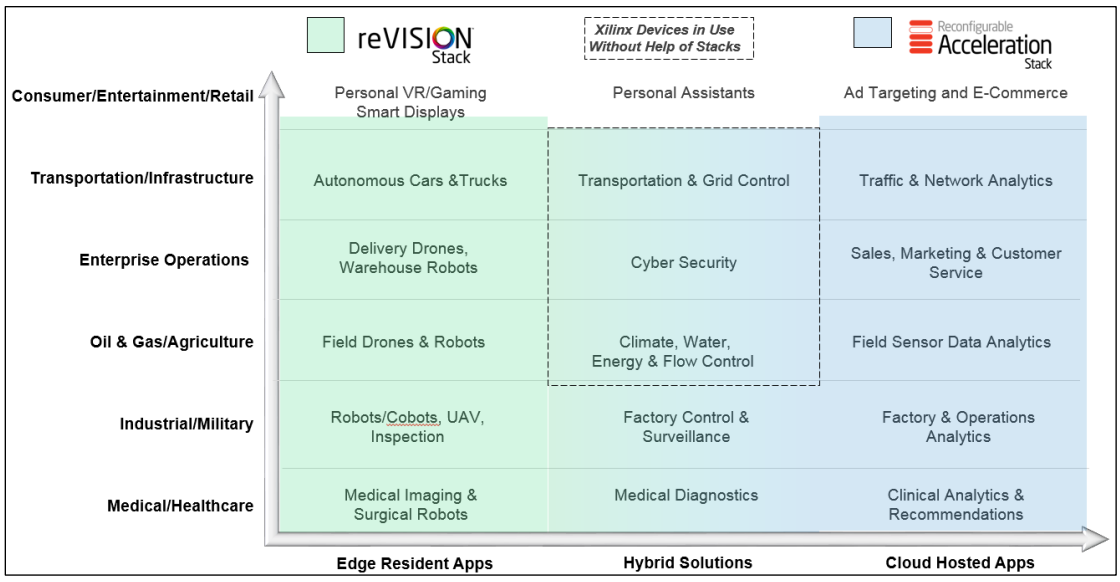


***Figure 1: Xilinx Broadening the Deployment of Machine Learning Applications from the Edge to the Cloud***
***Source: Machine Learning Landscape - Moor Insights & Strategy Research Paper***

## XILINX IN VISION AND MACHINE LEARNING TODAY

As shown in Figure 2, Xilinx has enabled numerous companies to build advanced embedded vision systems, deployed by 23 automotive manufacturers in 85 different car models for ADAS, and by 100's of additional embedded vision customers in 1000's of applications. At least 40 of these companies are already developing or deploying machine learning technologies to dramatically increase systems intelligence. Today, most of Xilinx's vision customers include engineers with significant hardware expertise, primarily targeting Zynq® All Programmable SoCs and MPSoCs.


*Figure 2: Xilinx Embedded Vision Momentum*

## reVISION TARGET APPLICATIONS AND MANDATES

Today, Xilinx is enabling applications in markets where differentiation is critical, systems must be extremely responsive, and the latest algorithms and sensors need to be quickly deployed. This includes applications in 'prosumer', automotive, industrial, medical and aerospace & defense markets, as well as high-end leading edge consumer. This typically excludes most high volume consumer and mainstream commoditized applications deploying less differentiated 'good enough' and mature technology.

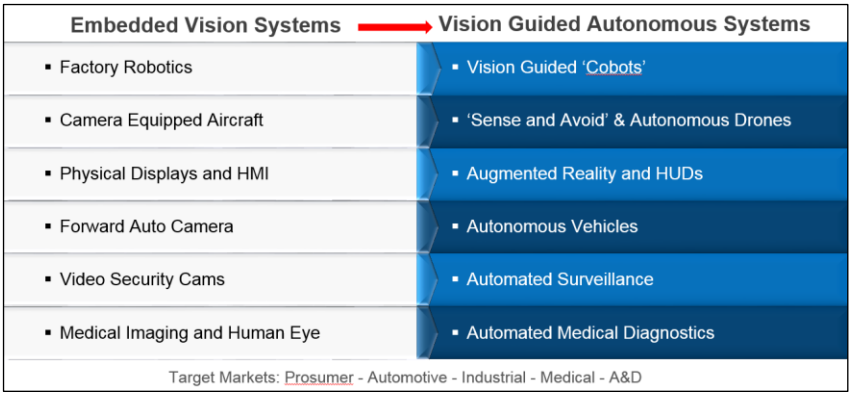


*Figure 3: From Embedded Vision to Vision Guided Autonomous Systems*

As shown in Figure 3, many of the traditional embedded vision applications are being dramatically transformed with the inclusion of machine learning and sensor fusion. Next generation applications include collaborative robots or 'cobots', 'sense and avoid' drones, augmented reality, autonomous vehicles, automated surveillance and medical diagnostics. These systems typically have three mandates:

1. Systems need to not only think, but immediately 'respond' to situations. This requires a more cohesive view from sensing to processing, analyzing, deciding, communicating and controlling. They also must be implemented very efficiently, deploying the latest machine learning inference techniques at 8-bit precision and below. Note that

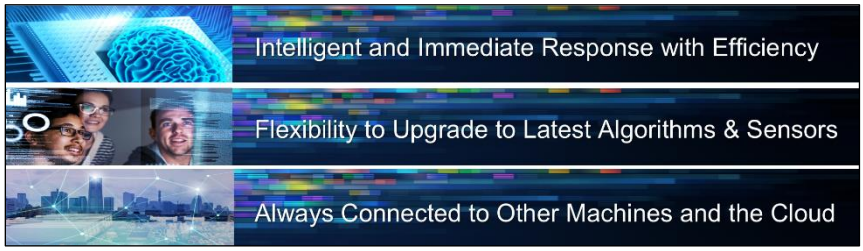


*Figure 4: The Application Mandates for Next Generation Vision Guided Systems*

technology optimized for machine learning training continues to diverge from technology optimized for machine learning inference. Xilinx has optimized its devices for inference.

2. Given the very rapid pace of change of neural networks and associated algorithms, and the rapid evolution of sensors, it is critical to enable the flexibility to upgrade systems through hardware/software reconfigurability.
3. As most of these new systems are now connected (IoT), they need to talk to legacy and existing machines, new machines that may be introduced in the future, and to the cloud. Xilinx defines this as 'any-to-any' connectivity.

Xilinx devices uniquely support each of these three mandates with significant and measureable advantages over alternatives. They enable the fastest response time from sensors through efficient inference and control, enable reconfigurability to support the latest neural networks, algorithms and sensors, and support any-to-any connectivity to legacy or new machines, networks and the cloud.


*Figure 5: Xilinx Unique Application Advantages*

While Xilinx devices today provide these advantages to expert users with hardware and RTL based design expertise, there has been a significant barrier to broader adoption; support for software defined programming using industry standard libraries and frameworks. The reVISION stack addresses this barrier to broad adoption.


*Figure 6: Barrier to Broader Adoption*

**reVISION STACK**

The Xilinx reVISION stack includes a broad range of development resources for platform, algorithm and application development. This includes support for the most popular neural networks including AlexNet, GoogLeNet, SqueezeNet, SSD, and FCN. Additionally, the stack provides library elements including pre-defined and optimized implementations for CNN network layers, required to build custom neural networks (DNN/CNN). This is complemented by a broad set of acceleration-ready OpenCV functions for computer vision processing. For application level development, Xilinx supports popular frameworks including Caffe for machine learning and OpenVX for computer vision (to

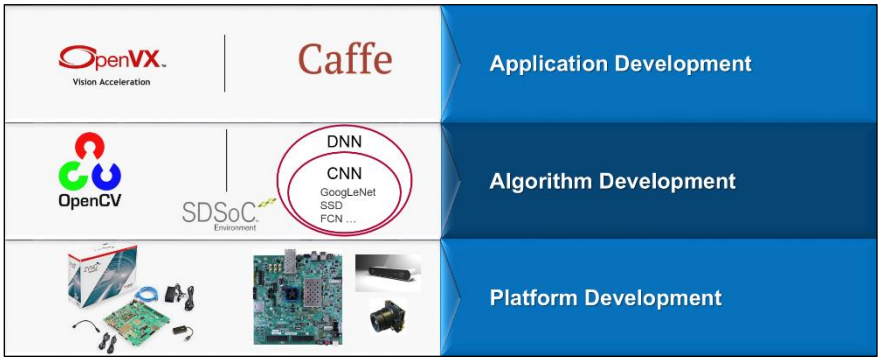


*Figure 7: Xilinx reVISION Stack*

be released later in 2017). The reVISION stack also includes development platforms from Xilinx and 3rd parties based on Zynq SoCs and MPSoCs.

## REMOVING THE BARRIER TO BROAD ADOPTION

The reVISION stack enables a broad set of design teams without deep hardware expertise to use a software defined development flow to combine efficient implementations of machine learning and computer vision algorithms into highly responsive systems.

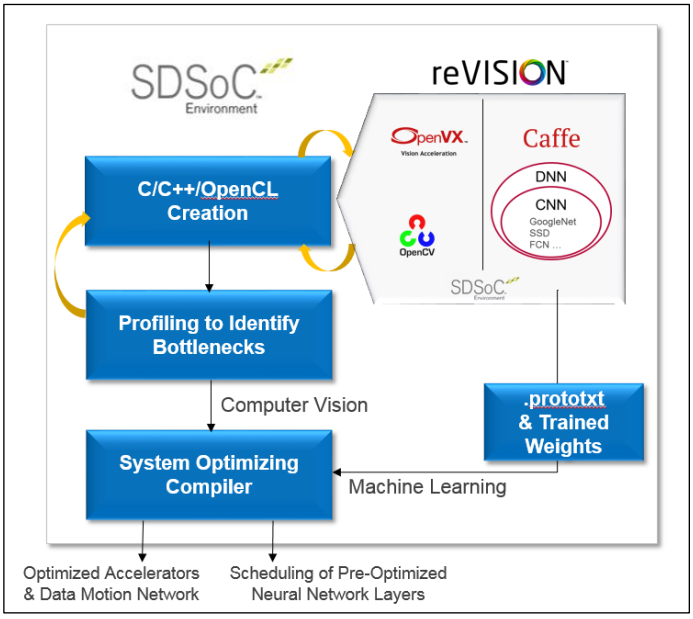


*Figure 8: The reVISION Software Defined Flow*

As shown in Figure 8, the reVISION development flow starts with a familiar, eclipse-based development environment using C, C++ and/or OpenCL languages and associated compiler technology; this is called the SDSoC development environment. Within the SDSoC environment, software and systems engineers can target reVISION hardware platforms, and draw from a pool of acceleration-ready computer vision libraries, and/or soon the OpenVX framework, to quickly build an application.

For machine learning, popular frameworks including Caffe are used to train a neural network, with the Caffe generated .prototxt file used to configure an ARM based software scheduler to drive the CNN inference accelerators pre-optimized for programmable logic.

For computer vision and other proprietary algorithms, users can profile their software code to identify bottlenecks and label specific functions in the code they want to speed-up and 'hardware optimize'. A 'system optimizing compiler' is then used to create an accelerated implementation, including the processor/accelerator interface (data movers) and software drivers. When combining computer vision and machine learning, this compiler will create an optimized fused implementation.

As shown on the left side of Figure 9, expert Xilinx users deploying traditional RTL-based design flows, working with ARM based software developers, spent considerable design time creating highly differentiated machine learning and computer vision applications.

To further speed design time and reduce the reliance on hardware experts, Xilinx introduced the SDSoC development environment approximately two years ago, based on C, C++ and OpenCL languages. While this significantly shortened the

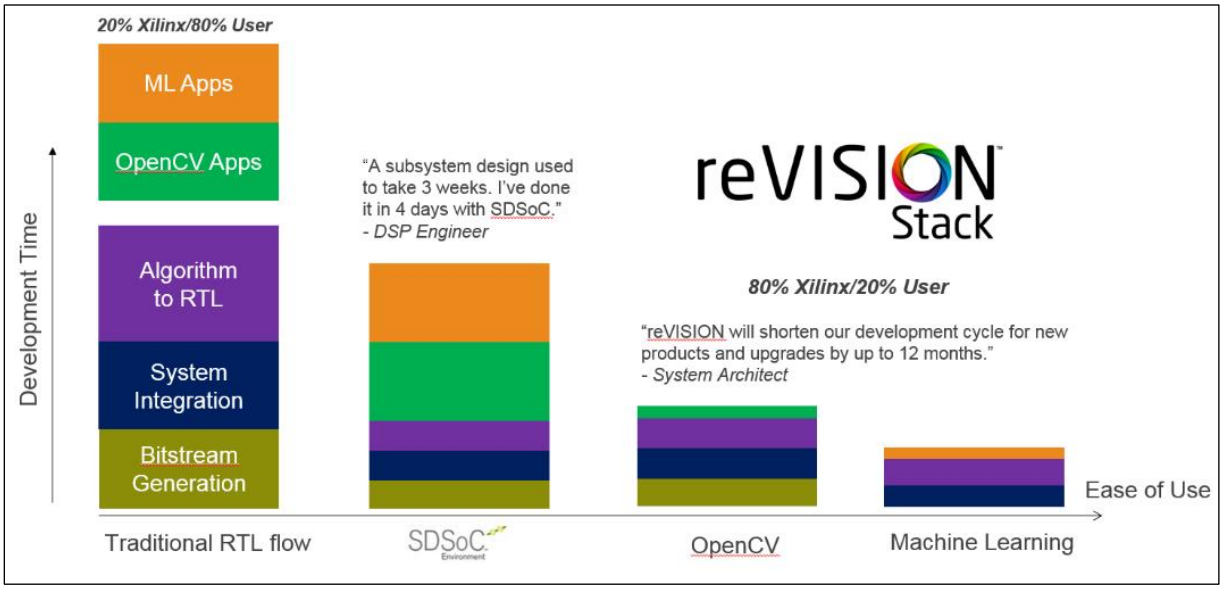


*Figure 9: reVISION Stack – Removing Barriers to Broad Adoption*

development cycle for an additional 1000 users who were capable of developing their own base platforms, libraries and applications, it still fell short of what was required for very broad adoption and deployment, and the complexities now being introduced by machine learning.

As shown on the right of Figure 9, Xilinx's new reVISION stack will enable a much broader set of software and systems engineers, with little or no hardware design expertise to develop intelligent embedded vision systems easier and faster, including the combination of extremely efficient implementation of machine learning and computer vision algorithms into highly responsive systems.

**RESPONSIVENESS – LOWEST LATENCY FROM SENSOR TO INFERENCE & CONTROL**

As stated above, the software defined reVISION flow enables the rapid development of highly responsive systems. In fact, when comparing performance metrics against embedded GPUs and typical SoCs, Xilinx significantly beats the best of this group, Nvidia.

Benchmarks of the Zynq SoC targeted reVISION flow against Nvidia Tegra X1 have shown up to 6x better images/second/watt in machine learning, 42x higher frames per second for computer vision processing, and 1/5 the latency (in milliseconds), critical for real time applications.

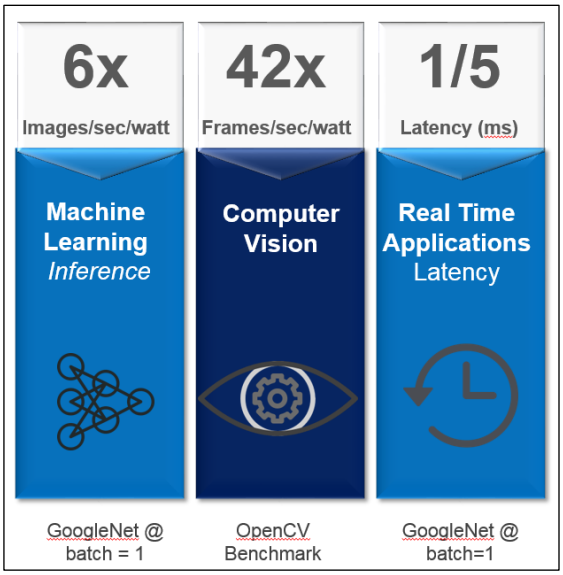


- Nvidia TX1 published: https://devblogs.nvidia.com/parallelforall/jetpack-doubles-jetson-tx1-deep-learning-inference/
- All Machine Learning Measurements taken at Batch = 1, ImageNet Images
- Computer Vision: Nvidia TX1 with 256 Cores; gpu::StereoBM_GPU; OpenCV 2.4.13
- Up to 6x: based on Xilinx roadmap

***Figure 10: Xilinx reVISION vs. Nvidia Tegra X1***

As shown in Figure 11, there is huge value to having a very rapid and deterministic system response time. This example shows a car with Xilinx reVISION based on Zynq SoCs relative to a car with an Nvidia Tegra device identifying a potential collision and deploying brakes. At 65 mph, depending on how the Nvidia device is implemented, the Xilinx response time advantage translates to a range of 5 to 33 feet of distance, which could easily become the difference between a safe stop and a collision.

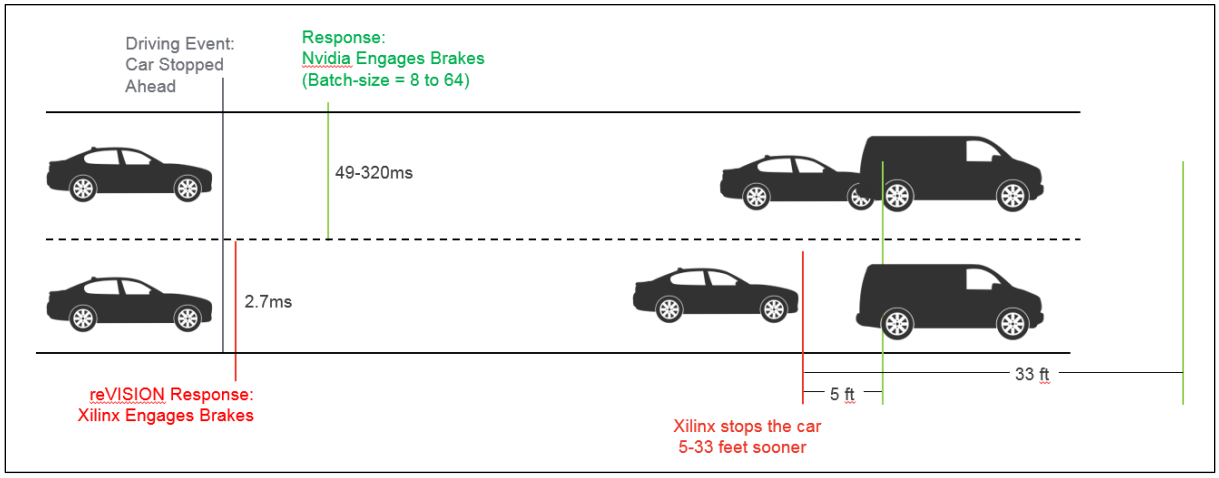


- Xilinx: ZU9 running GoogLeNet @ batch = 1
- Nvidia TX1: 256 Cores; running GoogLeNet @ batch = 8
- Assuming the car was driving at 65 mph

***Figure 11: Why Response Time Matters: Xilinx vs Nvidia Tegra X1***

These significant advantages in response time are derived from fundamental architectural advantages of Zynq SoCs relative to embedded GPUs and typical SoCs. As shown in Figure 12, embedded GPUs and typical SoCs require frequent access to external memory from sensors through vision, machine learning and control processing. In contrast, Zynq SoCs deploy an optimized and streamlined dataflow implemented with programmable logic and significantly more internal memory (up to 19x more than Nvidia Tegra X1). This not only results in 1/5 the latency relative to alternatives, but also deterministic response, critical for many real-time applications.

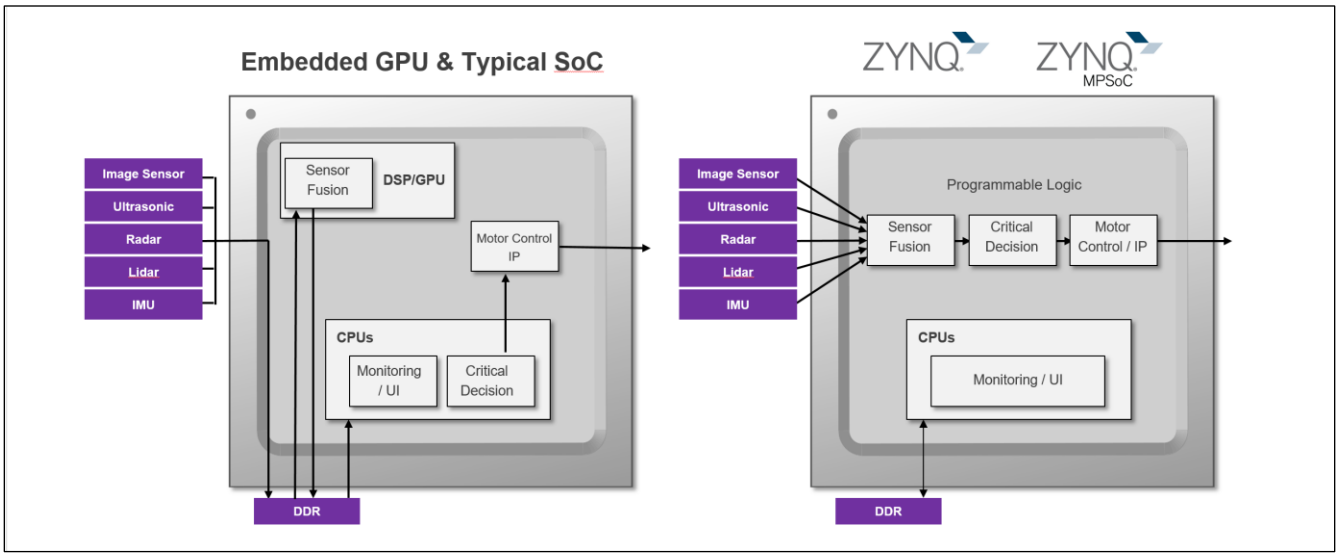


*Figure 12: The Source of Xilinx Response Time Advantages*

## RECONFIGURABLITY FOR LATEST NETWORKS AND SENSORS

While response time is important, Xilinx solutions also provide very unique advantages in reconfigurability. In order to deploy the best systems with the state of the art neural network and machine learning inference efficiency, engineers must be able to optimize both the software and hardware throughout the product lifecycle. As shown in Figure 13, the last two years of advancements in machine learning has generated more technology than the advancements over the last 45 years. Many new neural networks are emerging along with new techniques that make deployment much more efficient. Whatever is specified today, or implemented tomorrow, needs to be 'future proofed' through hardware reconfigurability. Only Xilinx All programmable devices offer this level of reconfigurability.

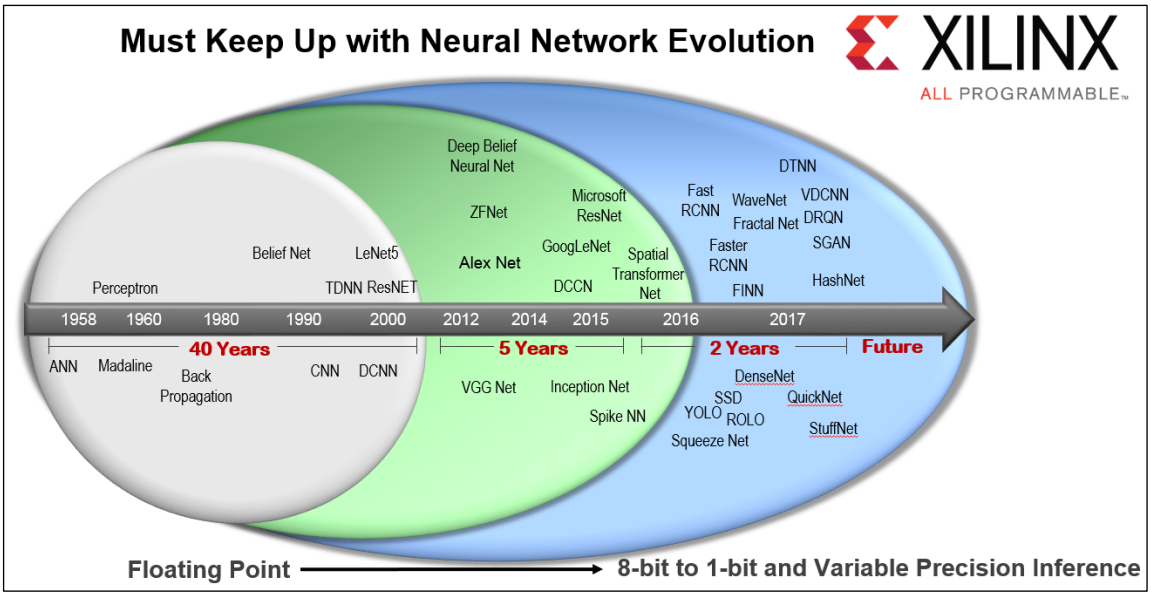


*Figure 13: Why Reconfigurability Matters as Machine Learning Evolves*

As shown in Figure 14, there is a similar requirement for reconfigurability to manage the rapid evolution of sensor technology. The AI revolution has accelerated the development and evolution of sensor technologies across numerous categories. It has also resulted in a mandate for a new level of sensor fusion, combining multiple types of sensors in different combinations to create a full and complete view of the system's environment and objects in that environment. As with machine learning, whatever sensor configuration is specified today, or implemented tomorrow, needs to be 'future proofed' through hardware reconfigurability. Again, only Xilinx All programmable devices offer this level of reconfigurability.

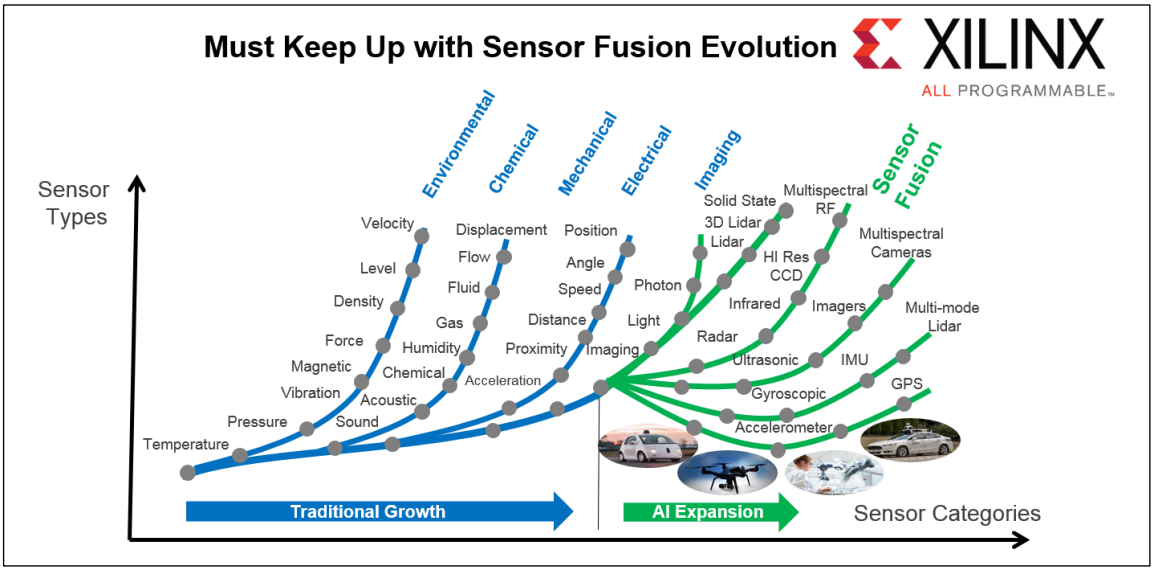


*Figure 14: Why Reconfigurability Matters as Sensors Evolve*

## ANY-TO-ANY CONNECTIVITY AND SENSOR INTERFACES

As shown in the Figure 15, Zynq based vision platforms offer robust any-to-any connectivity and sensor interfaces.

Zynq sensor and connectivity advantages include:

1. Up to 12x more bandwidth relative to alternative SOCs currently in the market, including support for native 8K and custom resolutions.
2. Significantly more high and low bandwidth sensor interfaces and channels, enabling highly differentiated combinations of sensors such as RADAR, LiDAR, accelerometers and force torque sensors.
3. Industry leading support for the latest data transfer and storage interfaces, easily reconfigured for future standards.

| | Movidius Myriad2 | Ambarella A12S, S3L | NXP Freescale i.MX8 | Qualcomm Snapdragon 8xx | Renesas H3 | NVIDIA Tegra X1 | Xilinx ZUEV7 | Advantage |
|---|---|---|---|---|---|---|---|---|
| MIPI Interfaces / Camera Support | 12 MIPI lanes 6x HD cameras | 8 MIPI lanes 2x 4K cameras | 8 MIPI lanes 2x 4K cameras | 8 MIPI lanes 2x 4K cameras | 16 MIPI lanes 8x HD cameras | 12 MIPI lanes 6x HD (2x 4K) cameras | 96 MIPI lanes 18 - 48 cameras w/ up to 8K | **12x more bandwidth Only 8K option Custom Resolutions** |
| Video Interfaces | MIPI DSI | HDMI 1.4, HD-SDI | HDMI, DSI | HDMI 2.0 | 3x low res HUDs | HDMI 2.0 (Output) | HDMI 2.0 (Input/Output), DisplayPort 1.2/1.4, 12G-SDI, MIPI-DSI | **Multitude of 4K interfaces combinable to support 8K and custom resolutions** |
| High Bandwidth Sensors: RADAR, LiDAR… | 1x 1GbE | 1x 1GbE, SPI | 2x 1GbE (AVB), 3x CAN-FD, 4x SPI, | | 1x CAN, 1x 1GbE (AVB), SPI | SPI | 48x: CAN/CAN-FD, 1GbE (AVB), SPI | **More Smart Sensor interfaces than any other SOC (up to 48 channels)** |
| Low Bandwidth Sensors: Accelerometer, Force torque… | 3x I2C, 1x UART | 1x I2C, 1x, 1x UART | 5x I2C, 5x UART | | I2C | | I2C, UART, GPIO | **More channels of lower bandwidth IO than any other SOC** |
| Data Transfer & Storage Interfaces | USB 3.0, 2x SDIO (host), USB 3.0 | USB 2.0, eMMC, NAND/NOR, SDIO, SDXC | PCIe x2 | SD, eMMC, SATA, USB 2.0, 3.0 | USB 2.0, 3.0, SATA, SD | USB 3.0, SDIO, eMMC | USB 2.0, 3.0, PCIe Gen 1.0/2.0/4.0 (PL) x4, x8, 10GE, SATA 3.1, NAND/NOR, SD/eMMC | **Industry leading support for *high bandwidth* data transfer and storage interfaces** |
| Based on publicly available sources | | | | | | | **All with support for news standards and/or new sensor technologies** | |

***Figure 15: Xilinx vs. Nvidia and SoCs in Sensors/Connectivity***

## XILINX vs. INDUSTRY ALTERNATIVES

By combining the unique advantages of Zynq-based platforms with software defined development environments equipped with libraries and support for industry-standard frameworks, reVISION is positioned as the best alternative for vision-based system development. As discussed, reVISION uniquely addresses all 3 'mandates' for intelligent applications where differentiation and time to market with the latest technologies is critical; Responsiveness, Reconfigurability, Any-to-Any Connectivity. It also removes the barrier to broader adoption with software defined programming.

As shown in Figure 16, on the vertical axis, only reVISION can enable optimization from sensors through machine learning inference and connected control, yielding the best system response time. On the horizontal axis, only reVISION can provide the level of reconfigurability for hardware optimized acceleration of algorithms and upgrades to the latest sensors and connectivity needs. While Xilinx devices have enabled many 100's of

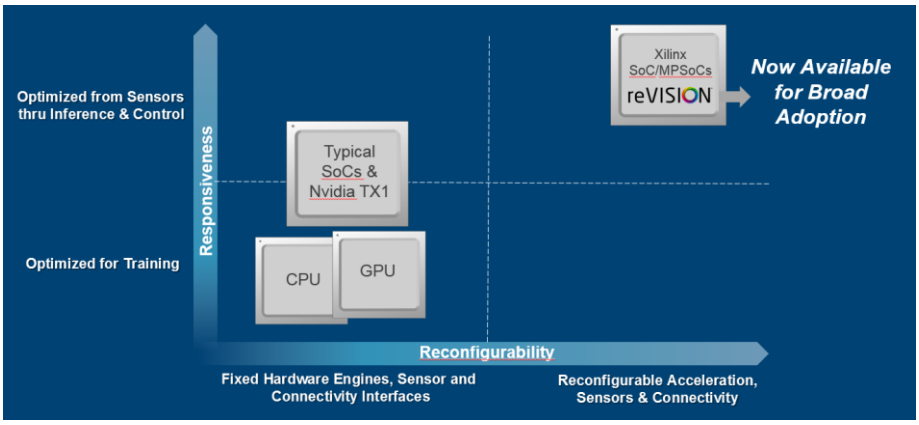


***Figure 16: Xilinx – Most Responsive and Reconfigurable Alternative in Target Markets***

customers with hardware experts to realize these advantages in the past, the new reVISION stack removes the barrier to broad adoption by enabling software defined programming using industry standard libraries and frameworks.

**IN SUMMARY, reVISION:**

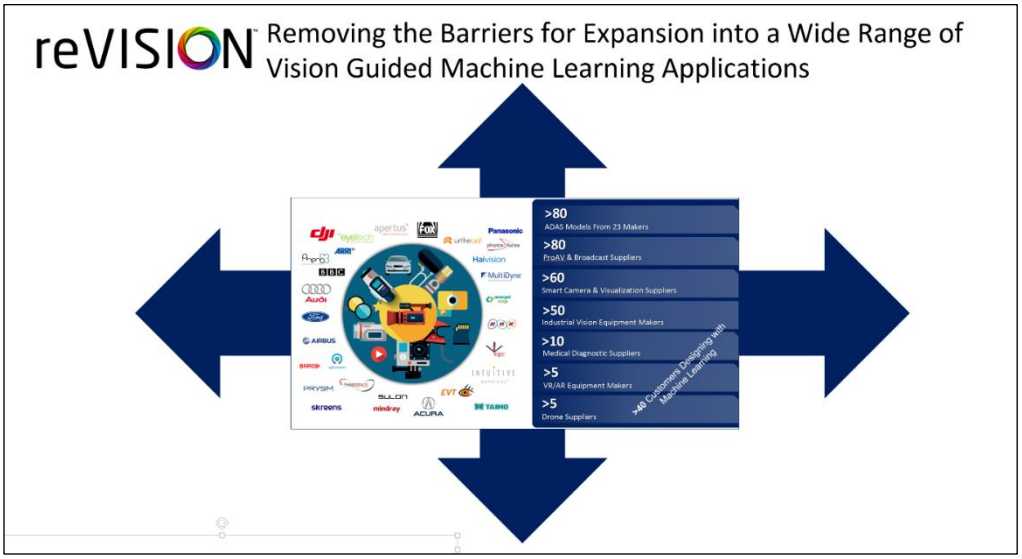


Xilinx is expanding into a wide range of vision guided machine learning applications with the Xilinx® reVISION™ Stack. This announcement complements the recent Reconfigurable Acceleration Stack, significantly broadening the deployment of machine learning applications with Xilinx technology from the edge to the cloud. The new reVISION stack enables a much broader set of software and systems engineers, with little or no hardware design expertise to develop intelligent vision guided systems easier and faster. These engineers can now realize significant advantages when combining machine learning, computer vision, sensor fusion, and connectivity.