INUITIVE INTRODUCES THE NU4000, AN ADVANCED 3D IMAGING, DEEP LEARNING AND COMPUTER VISION PROCESSOR - THE ULTIMATE SOLUTION TO 'XR' (VR/AR/MR) DEVICES DESIGN CHALLEGES

Inuitive, a developer of cutting edge 3D Imaging and Computer Vision and image processors, has introduced the NU4000, an advanced multi-core vision processor that supports 3D Imaging, Deep Learning for Artificial Intelligence and Computer Vision processing for Augmented Reality, Virtual Reality, Drones, Robots, and many other applications.

The NU4000 is the most advanced 3D Imaging and Vision processor to date. It was designed specifically with XR use cases challenges in mind. Before outlining the NU4000 unique benefits and capabilities, let's examine the XR use case challenges.

XR (VR/AR/MR) emerging use cases deliver exciting application possibilities to the market, while representing significant design challenges to device developers and manufacturers. Below are six common and key challenges with the XR market:

- Multiple sensors with loads of pixels to process, along with tight synchronization and fusion requirements. An advanced XR device may integrate up to 6 to 10 CMOS sensors that collectively generate huge pixels throughput, introducing substantial challenges to synchronize these sensors in order to enable fusion of them efficiently.
- Such loads of image data exist in order for them to be processed. This translates into a need for heavy computer vision processing power that is difficult to achieve. Furthermore, extreme memory throughput demand typically leads to a situation whereby memory bandwidth, rather than compute power, transpires as the limiting factor of the system performance.
- XR is aimed at addressing challenging use cases, such as gaming, in which low latency is a critical factor in achieving an optimal user experience. Specifically, the Motion-to-Photon latency is one of the most important system parameters influencing the UX. Minimizing the latency of imaging and vision processing tasks to a timeframe of a few milliseconds is one of the major challenges in system design.
- Performance vs. flexibility tradeoffs is certainly a key challenge. While utilizing fully-programmable compute cores provides the desired flexibility and future proof, it definitely transpires into power inefficiency, while simultaneously insufficient for addressing a variety of the above-mentioned challenges for latency minimization.
- Ultimately, mobile devices are a dominant player in the XR device domain, and those that are not will eventually become mobile and battery powered. Hence, power consumption plays a vital role in these particular use case challenges, and addressing

the above-mentioned compute requirements, while maintaining reasonable power consumption levels, is a major challenge.

In the last 3 years, Inuitive has actively invested effort and time in the development and integration of XR devices. This initiative stems from various business engagements that revolve around our first generation 3D Imaging & Vision Processor – the NU3000. In collaborating with partners and customers, we inevitably faced and learned about the aforementioned challenges. Based on thorough architectural research, we are able to incorporate several key capabilities into the NU4000 processor to make it the ultimate vision processing solution for XR devices.

Let us examine key architectural elements of the NU4000 processor that powers its leveraged efficiency for XR use cases:

Adoption of optimal computing technology for diverse computational, system tasks

At Inuitive, we believe that in order to generate a highly effective computing solution, it's vital to employ the most efficient processing technology for each specific computing task.

Of the various examples, the following are two noteworthy scenarios in which this principle is applicable:

The technology allows you to run heavy computing tasks, such as stereo depth matching by DSP or GPU, which is simpler than developing an efficient, dedicated HW engine. However, this and other like approaches will yield limited throughput and extreme power consumption. Thus, Inuitive has implemented a proprietary, highly-efficient, high-throughput HW depth processing engine that supports both Stereo and Structured Light technologies. The same notion is applicable for Artificial-Inelegance/Deep-Learning tasks. It's possible to run CNN (Convolution Neural Networks) computing by CPU, GPU or VPU. However, it is impossible to achieve the market required CNN computing power at desired energy efficiency levels without utilizing a dedicated CNN processor that can process hundreds of MACs (Multiply Accumulate) operations per a single computing cycle. For this purpose, we integrated such a dedicated CNN processor based on Synopsys EV6x architecture. The following diagram describes the main architectural blocks of NU4000, illustrating the utilization of different computing architectures for a range of tasks:

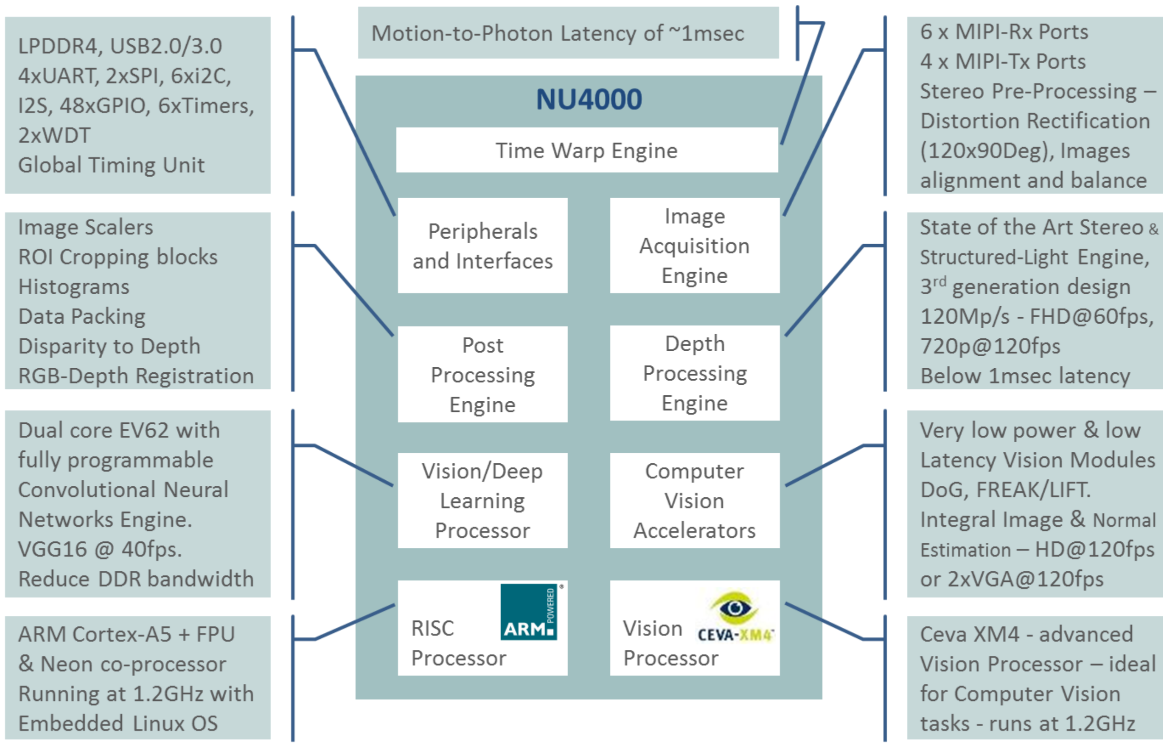

Careful balance between hard-wired processing blocks and full-programmable ones:

As implied above, it can be very tempting to build an architecture that provides a few Terra OPS (Operations Per Second) by integrating multiple programmable cores, and then leave it for the troubled, embedded SW developers to cope with the challenge of integration and optimization of multiple heavy computing tasks that continue to pile up in the production processes of one device model to the next. It seems like a great solution that provides extensive flexibility required, and such solutions do occasionally present themselves in the market. However, this is neither a practical, nor effective means to approach coping with the above listed challenges. In order to deliver all of the required computing capabilities at the desired power consumption level, it is essential to intelligently combine dedicated HW accelerators with fully-programmable cores. This is precisely what the team at Inuitive does best, as NU4000 introduces a novel approach to develop Computer Vision processing blocks that are based on HW pipeline implementation. The NU4000 SLAM accelerator unit (Computer Vision Accelerators), for example, enables highly accurate feature-point extraction by implementing heavy algorithms, such as DoG and FREAK, via a dedicated HW pipeline that enables extraction of hundreds of key-point descriptors at 120fps from a dual HD tracking camera structure.

Dedicated imaging and vision pipelines releasing memory access bottlenecks

Another challenge that continues to grow in scope in adopting multi-core architecture is the need for extreme memory access bandwidth. An issue of this nature typically extends when computing cores are processing data stored either on on-chip SRAM banks, or on external DDR memory. The architecture that we developed at Inuitive enables the routing of image data directly from the cameras to the main processing pipelines. This applies to both the depth processing pipeline as well as the SLAM related accelerators. Such an approach significantly reduces required memory bandwidth, (and space), since the high throughput cameras data is not directed towards memory, rather solely the compact meta-data, (which is the output of the processing pipelines), is directed towards the memory blocks for further processing. This approach is illustrated in the following diagram:

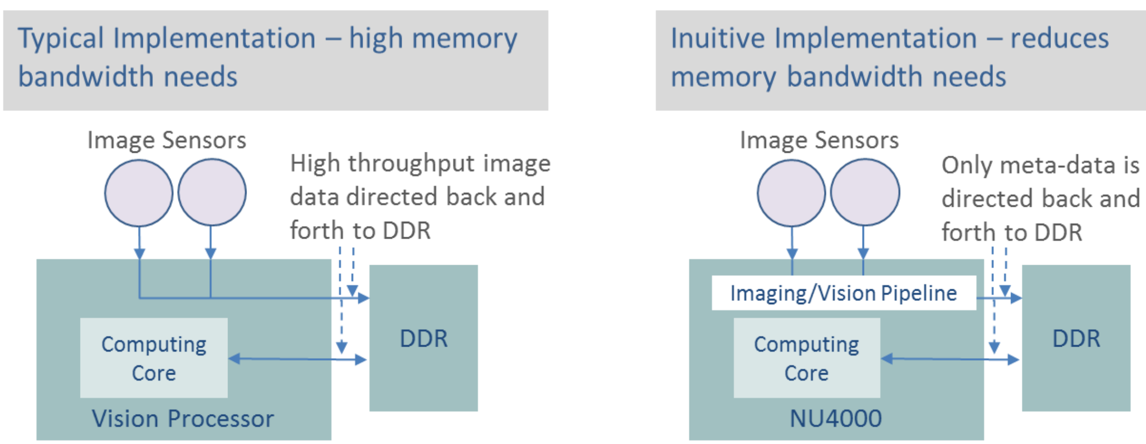

Unique architectural approach for reducing Motion-to-Photon latency

In XR use cases, UX is strongly impacted by "Motion-to-Photon" latency, hence it is critical to minimize the time between detection of a user's head pose movement and the appropriate change of the image projected on the display. NU4000 introduces an advanced image re-projection implementation, (AKA "Time Warp" implementation), minimizing motion-to-photon latency to up to 1 msec. This solution is based on 3 main elements:

1. Fast and accurate 6DoF tracking engine on NU4000 that utilizes SLAM accelerators for precise detection and calculation of a user's head movements.
2. Dedicated HW pipeline, enabling image re-projection on-the-fly, (as well as image warping and chromatic aberrations correction), asynchronously projects onto the display data stream.
3. Fast "loop-closure" between the timeframes of 6DoF tracking info and the image re-projection HW engine.

Encapsulating the above 3 capabilities on the same chip enables the option to perform very fast orientation correction of the displayed images without the need to relay it on the relatively long rendering pipeline of the GPU. In addition, the need to relay it on the application processor is eliminated, thus bypassing the Motion-to-Photon latency that such a pipeline introduces. This is a novel system level solution that relays onto the unique architecture of NU4000, as illustrated in the following diagram:

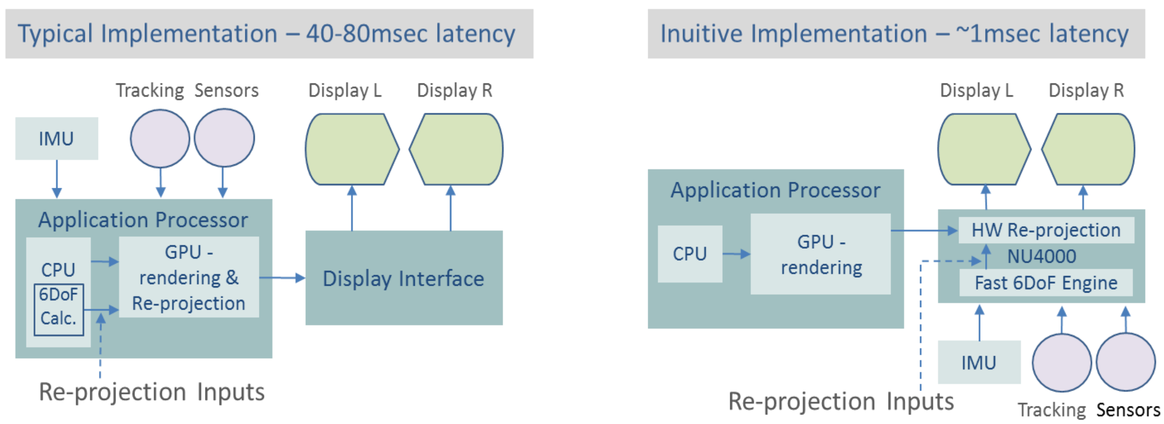

Plenty of fully-programmable vision compute power tightly coupled with system sensors

Efficient low-latency imaging pipelines are of great value, but are typically dedicated to certain limited tasks they perform very effectively. Nevertheless, we are actively a part of a highly dynamic market in which new applications and use cases constantly present themselves. In order to cope with such an environment, an effective vision processor must supply sufficient fully-programmable cores partnered with the right SoC architecture infrastructure, and a set of high-productivity tools that enable developer utilization of such available computing power. Inuitive certainly offers such compute power on NU4000, with appropriate SoC architecture enabling the multiple cores to access shared memory on the chip, sufficiently immense to store significant parts of the processes image data.

NU4000 provides more than 2.5 TOPS with programmable cores based on the following key elements:

- Synopsys EV62 dual-core vision processor cluster with CNN3.0 co-processor running at 1.2GHz
- Ceva XM4 vision processor running at 1GHz
- ARM Cortex-A5 General Purpose Processor running at 1.2GHz with Linux OS execution environment

Combined, these cores provide 13,000 CoreMark CPU power, over 500 GOPS for generic Computer Vision processing, and more than 2 TOPS for Neural Networks processing.

Dedicated HW mechanism for sensor synchronization and fusion

System and SW developers in our domain often share their challenges with us about achieving accurate time synchronization between multiple sensors in the system and assigning precise time-stamps to them. This is definitely a challenge. With the Global Timing Unit that we implemented on NU4000, developer concerns regarding like issues are eliminated. SLAM implementation is simplified, and provides increased opportunity to focus on improving its vision algorithms, rather than spending time on achieving desired timing accuracy.

Architecture optimized for energy efficiency

Offering such extensive compute power on one device may lead to a very power-hungry chip that limits system battery time. A set of architectural concepts have been implemented on NU4000, allowing Inuitive to provide desired computing performance, while maintaining decent power figures, both at the chip and system levels. A variety of such concepts include:

- Using dedicated HW Accelerators for frequently used heavy computing tasks – an order of magnitude power consumption reduction
- Utilizing Vector Processors for Computer Vision tasks – the most power efficient architecture for generic CV algorithms execution
- Deep-Learning/Artificial-Intelligence tasks execution by dedicated optimized processor – x10 less power than GPU implementation
- Tight clocking management scheme that optimally minimizes the activity factors of different computing blocks
- Utilizing advanced 12nm silicon fabrication process which provides ultimate power efficiency.
- Chip architecture is system-aware – implementing fast and accurate control of peripherals such as CMOS sensors that enables system power optimization by minimizing and sensors and illumination modules activity factor

As a summary of the above statements it can be said that when designing the NU4000, the Inuitive ASIC architects applied great emphasis on solving the design challenges of XR devices while leveraging on the company’s special experience in this field and close collaboration with leading market players. The result is a very powerful 3D Imaging and

Vision processor that provides the best balance between performance, flexibility, scalability, power consumption and cost.

About Inuitive

INUITIVE (www.inuitive-tech.com) optimizes consumer experiences and enhances competitive advantages in the areas of Augmented Reality and Virtual Reality, Drones, Robots and Autonomous Cars, to name just a few.
INUITIVE's intelligent offering combines, algorithms, ASIC and System solution to realize the AI practice enabling devices to acquire more human capabilities. With AI at its core, The INUITIVE platform also includes a 3D Depth Sensing Computer Vision processor and powerful deep learning capabilities to enable smart devices to become even smarter.