

NATvision



NATvision is a complete environment for the development and deployment of sophisticated image and video processing applications. The environment consists of a broad range of hardware products combined with selected software technologies and engineering support from N.A.T..

Designed as a machine vision processing platform, NATvision offers higher performance and lower development costs than comparable solutions, with unlimited scalability and number of cameras.

Applications include automated visual inspection in manufacturing or transportation, medical image analysis and public safety camera aggregation.

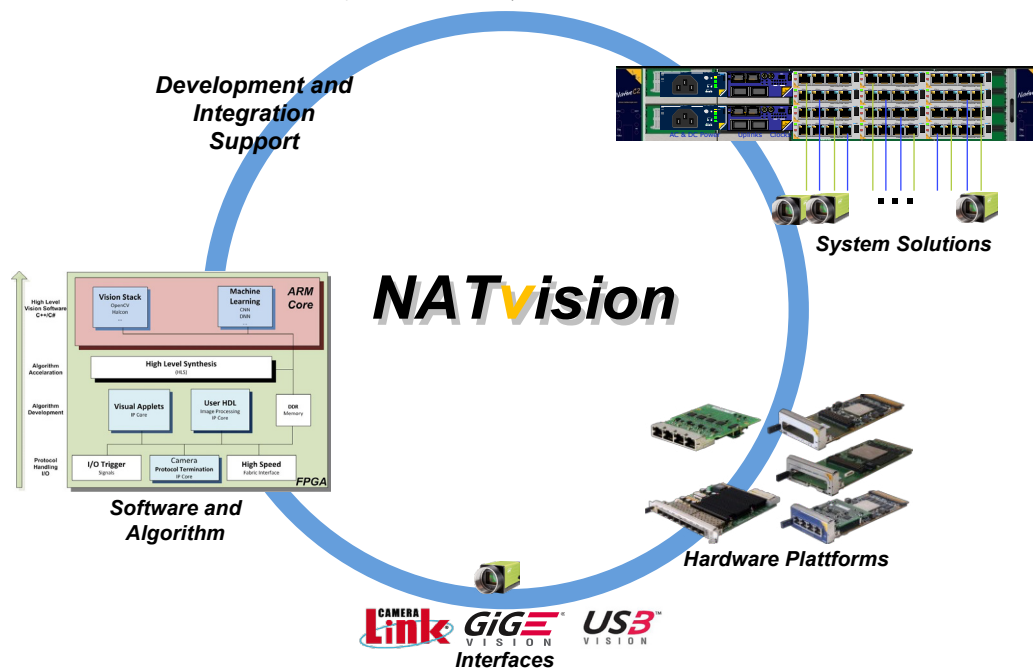
Learn more in this white paper.

NATvision is a complete environment for the development and deployment of sophisticated image and video processing applications. The environment consists of a broad range of hardware products combined with selected software technologies and engineering support from N.A.T..

Designed as a machine vision processing platform, NATvision offers higher performance and lower development costs than comparable solutions, with unlimited scalability and number of cameras.

The NATvision technology uses today's state-of-the-art FPGA resource boards, combining the performance and real-time advantages of FPGA based algorithms with software-based image processing.

NATvision systems are based on MicroTCA (MTCA), an open system standard for embedded applications, combining flexibility and scalability with high speed data rate capabilities. However, NATvision components can also be used standalone or be integrated into any custom hardware.



Applications

NATvision targets any application requiring high speed video or image processing or real-time manipulation, such as:

Vision inspection

- Quality inspection, failure detection
- PCB assembly inspection
- Railway inspection
- Machine vision

Analysis and merging

- Medical picture analysis and merging

Video – public safety

- Camera link aggregation
- Filtering and analysis
- Compression and re-coding



Vision technology based on MicroTCA

Compared to other system standards MTCA provides an unbeatable freedom of scalability and flexibility for vision applications ranging from single camera systems up to multiple camera systems combined with the processing power of FPGA and/or CPU/GPU units.

Scalability

MicroTCA allows you to build small and cost-effective systems with only one or a few cameras, up to high performance systems with support for up to 48 cameras in a single 2U high chassis.

Flexibility


The modular nature of MicroTCA enables you to combine camera systems with image processing units such as FPGA/CPU/GPU/DSP boards, allowing you to build system solutions which suit the needs of your application and meet your exact requirements.

High speed data path

The backplane-based bus topology of MTCA provides high speed interconnect links to exchange video or image data between computing boards or to combine multiple video streams into one data stream with up to 40Gb performance per slot. For applications requiring high-availability, MTCA can also provide redundancy, meeting system fail-over scenarios using a dual-star topology.

Expandability

Multiple MTCA systems can easily interchange data, clock and timing signals and distribute these over long distances. Hierarchical, pipelined and other system architectures can be configured. The number of connected and synchronized cameras increases with each MTCA system connected to the architecture.




Hardware building blocks

NATvision comes with a broad range of hardware products providing FPGA / CPU / DSP resources for image and video pre-processing, such as:

- Intel® ARRIA10 FPGA/ARM carrier for FMC modules
- Xilinx® Kintex® 7 FPGA/ARM carrier for FMC modules
- Xilinx Ultrascale+® FPGA/ARM for FMC modules (under development)
- Xilinx Kintex 7 FPGA module (double size) with 8 SFP interfaces
- FMC Interface module with 4 GigE Interfaces and Power-over-Ethernet (PoE)

Optional image and video post-processing solutions can be provided by:

- DSP modules with 8 multicore DSP chips
- PowerPC® processing modules with up to 8 cores and data path acceleration
- Third party Intel processing modules offering up to 4 cores running desktop Windows® or Linux
- Third party GPGPU modules with 2 NVIDIA® Tegra-K1® GPUs
-



Supported video/camera interface standards

NATvision support the following camera interface standards:

- GigE Vision®
- Camera Link®
- USB-Vision®



NATvision software architecture

The NATvision software architecture combines the advantages of FPGA based algorithm acceleration with the flexibility of on-chip CPU power. The picture below shows the software packages / technologies incorporated into the NATvision environment.

- Visual Applets: Block level design eliminates any FPGA knowledge
- OpenCV and Halcon image processing libraries with special adaptations for FPGA acceleration
- ARM® C/C++ code acceleration with High Level Design Synthesis (HLS) for FPGA acceleration
- Verilog or VHDL design entry for FPGA as example project
- Linux BSP as software for the ARM® CPU

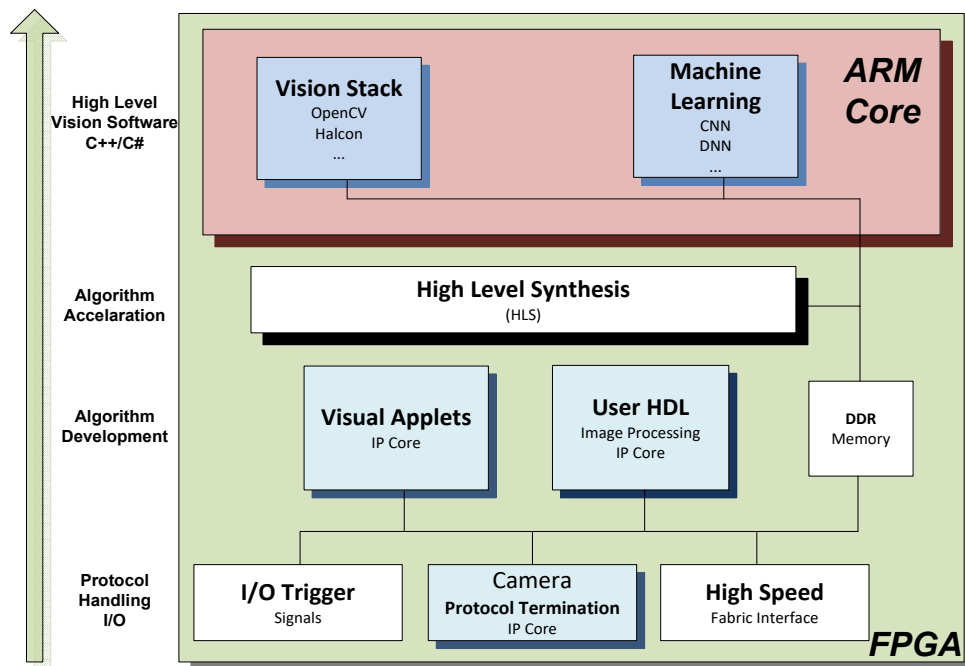


Figure 1: NATvision Software Architecture

Image processing libraries: OpenCV and Halcon

OpenCV (Open Computer Vision) is an open source library offering functions for machine vision, which was initially developed by Intel. It is available for C, C++, Python and Java applications. Apart from its performance, the big advantage of OpenCV is its free availability and the richness of algorithms based on latest research results. OpenCV runs on several platforms such as ARM or x86.

A subset of OpenCV functions includes:

- Face tracking
- Machine learning
- Motion detection
- Object detection
- Common image processing filters: Sobel, Canny, Gauß

The machine vision library HALCON by MVTec Software GmbH is focused on industrial image processing applications, a key part of 'Industry 4.0'. It has a feature-rich image processing library with over 2000 operators plus an integrated development environment (IDE, called HDevelop). It is available for C, C++, C# and Visual Basic and runs on several platforms such as ARM or x86 and uses GPU resources for acceleration.

A subset of HALCON functions includes:

- 3D machine vision
- Deep-learning
- Object detection (matching)
- Detection of letters and fonts
- Detection of lines, circles and ellipses

Graphical design with Silicon Software - VisualApplets

VisualApplets is a graphical design package for the design of image processing algorithms developed by Silicon Software.

The advantage of VisualApplets is that the hardware infrastructure is abstracted and hidden from the user. Thus, any knowledge about Hardware Description Language (HDL) or FPGAs is not required, allowing the user to focus on the image processing algorithm. Algorithms will finally compile into an IP black box which is part of a pre-synthesized FPGA design.

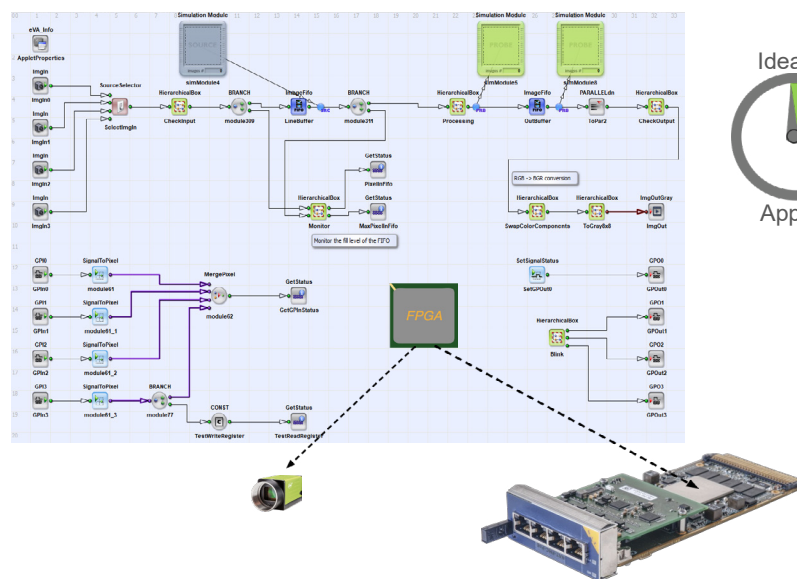


Figure 2: VisualApplets, idea to applet

Flow chart based design

In VisualApplets an image processing solution is developed in the form of a graphical flow chart in which operators are picked out of a comprehensive image processing library. It is also possible to verify the image processing system during development using simulation blocks that can be used for image generation and output visualization. This technique of visual debugging significantly decreases the development time as it is not necessary to do a compilation of the FPGA design. Figure 3 shows how to build an example processing chart that takes out the Y component of an incoming YCbCr data stream and forward the image output as grayscale RGB.

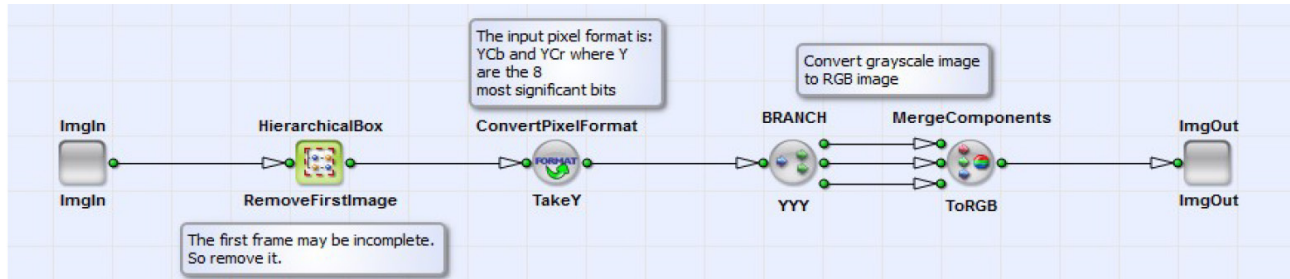


Figure 3: VisualApplets [Silicon Software]

Embedded VisualApplets - Example platform implementation

Figure 4 shows a typical FPGA processing system containing the embedded VisualApplets IP black box core implemented on a N.A.T. FPGA board. This infrastructure is pre-generated once by N.A.T., allowing customers to select the FPGA processing platform from the VisualApplets host PC application. Although the infrastructure containing the vendor IP blocks is completely hidden, the user can adjust several parameters and interfaces specific to their demands, such as the input and output. This allows the user to switch between different cameras and fabric distribution interfaces without touching the FPGA processing infrastructure.

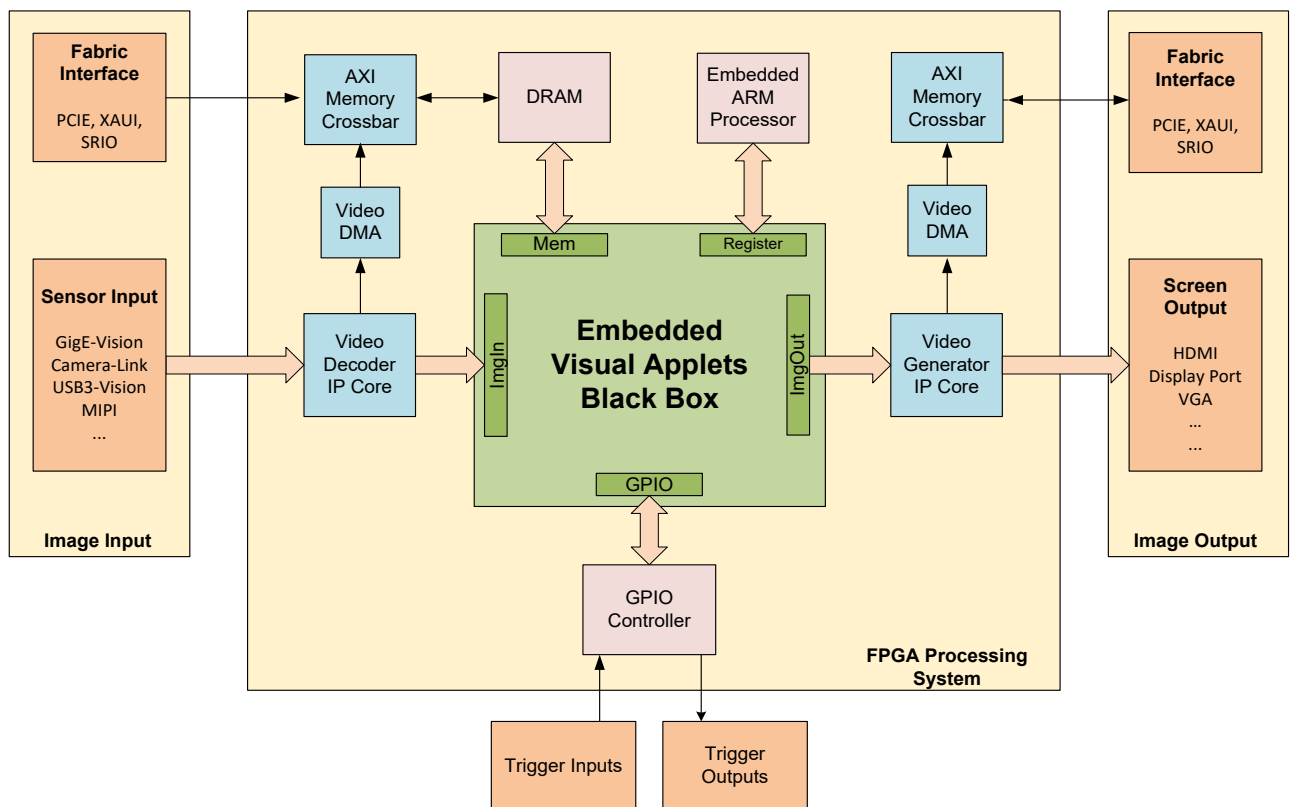


Figure 4: eVA black box IP implemented in FPGA processing system

Xilinx FPGA development environment

Xilinx System on Chip (SoC) FPGAs combine an ARM processor with an FPGA on the same chip. While the ARM processor is used to configure the processing platform and execute software-based algorithms, the FPGA can focus on video processing which is time critical and consumes performance. By using parallel pipeline streaming, the FPGA is more effective in video processing than a CPU. The Xilinx Vivado® software platform provides a graphical programming environment which enables users to build up a system with prebuilt Intellectual Property (IP) blocks for the FPGA. The Vivado library already contains a multitude of IP blocks, which are ready-to-use for building a video processing system. These blocks can be connected to build a typical image processing architecture as shown in Figure 5.

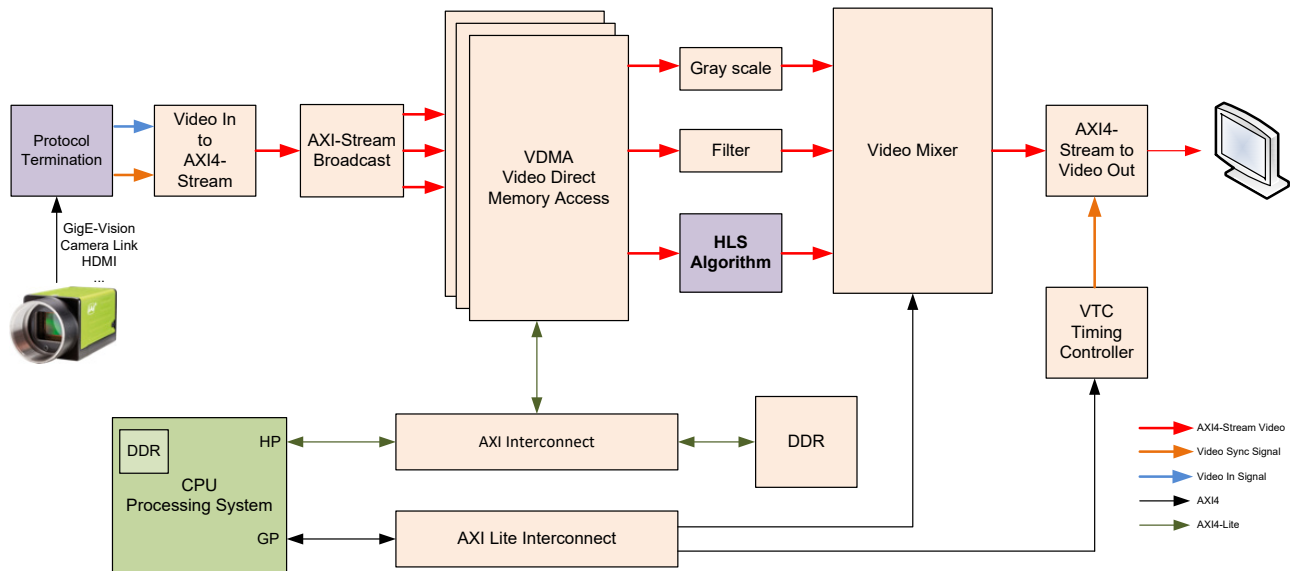


Figure 5: Xilinx image processing platform

Video input

Usually video or image signals from a camera sensor use a video streaming format that transports pixel data serialized line by line, combined with timing information that is needed to reconstruct the line data back to a 2D picture or video. To use industrial standard protocols such as GigE Vision it is required to do a protocol termination. The N.A.T. GigE Vision IP core provides this functionality.

Image processing algorithms

There are different sources available for image processing algorithms. Several image processing cores are available in the design tools IP catalogue, that provide basic conversion functionalities:

- Color space conversion
- Chromatic resampling
- Image scaling

In order to access more complex functions, developers can use high level synthesis tools such as Vivado HLS.

High Level Synthesis (HLS)

High Level Synthesis allows a developer to write FPGA code using C, C++, or System C and C- based libraries such as OpenCV or Halcon and have it generate the HDL in the form of VHDL. It is also more intuitive and less complex to describe image processing algorithms in a C-based language versus using VHDL. These loops are highly optimizable using FPGA resources. With the Vivado HLS platform it is possible, to synthesize C functions into VHDL and import it as an IP block into the video processing system.

Xilinx 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 (Convolutional Neural Network) network layers, required to build custom neural networks (DNN/CNN). The machine learning elements are complemented by a broad set of acceleration-ready OpenCV functions for computer vision processing. For application level development, Xilinx supports industry-standard frameworks including Caffe for machine learning and OpenVX for computer vision. The reVISION stack also includes development platforms from Xilinx and third parties, including various types of sensors. [Xilinx]

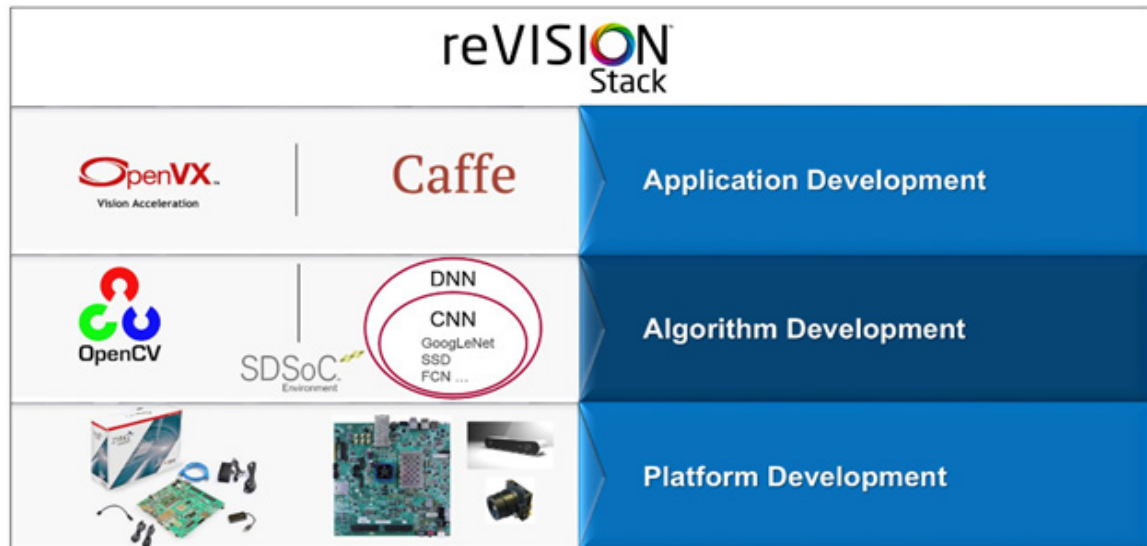


Figure 6: reVision Stack [Xilinx]

Off-Board data paths

A typical data flow chart using one of the N.A.T. FPGA boards is shown in Figure 7. Incoming camera data from up to 4 parallel GigE Vision cameras are interfaced using the FMC-4GigE-PoE module mounted on either the NAMC-ARRIA10-FMC, NAMC-ZYNQ-FMC or future boards with Ultrascale+ FPGAs. The picture shows how the embedded CPU interacts with the FPGA logic. In the GigE Vision use case, the embedded CPU configures the GigE Vision IP cameras using the control plane of the GigE Vision protocol termination IP block. The image data stream will be forwarded to the FPGA logic and processed by image processing user logic. The image data will then be stored on the internal memory or externally distributed using the backplane interfaces. While having a connection to the board memory the CPU can also be used for live data tracking.

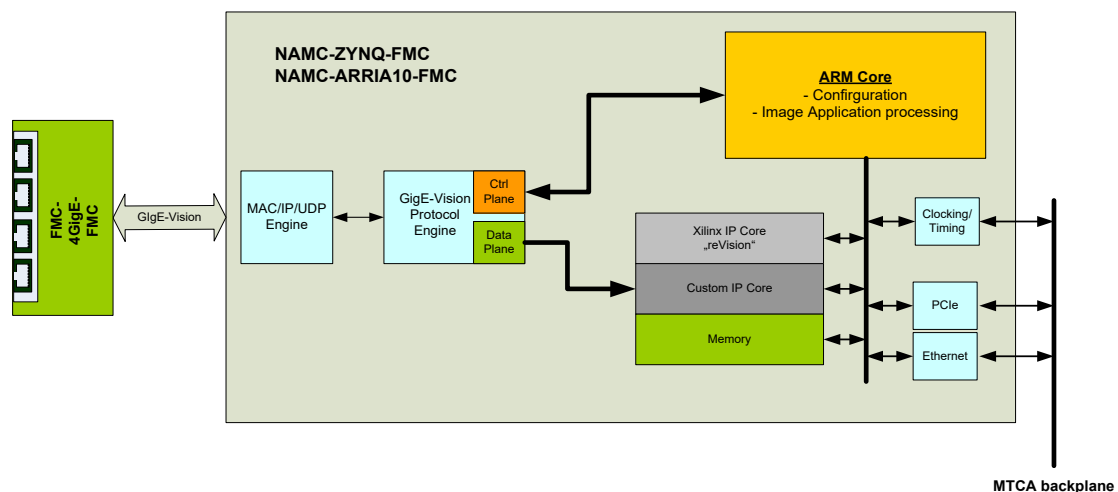


Figure 7: Off board data paths

FPGA cores by N.A.T

N.A.T. GigE Vision IP Core for Xilinx and Intel FPGAs

GigE Vision was developed in 2006 by the Automated Imaging Association (AIA) and uses the Ethernet protocol as transport media. GigE Vision supports multiple streams over very long distances using common Ethernet cables. As Ethernet is a widely used communication standard for industrial and consumer market, servers, personal computers and smart embedded devices can be used as receiver devices without requiring additional frame-grabber line interface cards. It is also possible to use the Power-Over-Ethernet (PoE) feature to use the camera with a single cable providing both image data and power.

The GigE Vision IP core by N.A.T. can be used for GigE Vision protocol termination on FPGAs for using compliant GigE Vision cameras. The IP core integrates the necessary UDP/IP infrastructure that is needed for communication with the camera, ensuring real time behavior, higher frame rates and resolutions.

Initial Release Date	May 2006		
Output configurations	Configuration	Image data throughput	Number of cables
	1 GigE	115 Mbytes/s	1 cable
	2.5 GigE	280 Mbytes/s	1 cable
	5 GigE	570 Mbytes/s	1 cable
	10 GigE	1100 Mbytes/s	1 cable
	WiFi	25 Mbytes/s	N/A
Camera control	Uplink channel	Symmetric with downlink channel	
	Downlink channel	Shared with image data	
	Trigger input signal	Hardware trigger on camera. Software trigger, optionally synchronized by Precision Time Protocol (IEEE1588)	
Receiver devices	Network interface card (NIC) can be on motherboard or inserted as an add-in card. Possibility to use a GigE Vision frame grabber.		
Supported transfer topologies	Point-to-point, multiple destinations	Direct connection to network card or to an Ethernet switch is possible. Support for multicast and broadcast.	
Cabling	Types	Max. length <i>(typical at 85 MHz)</i>	Power over cable <i>(wattage at camera)</i>
	CAT-5e/CAT-6a/CAT-7	100m	Optional 13W Optional 25W
	Multi-mode fiber	500m	No power
	Single-mode fiber	5000m	No power
	SFP+ direct attach	10m	No power

Figure 8: GigE Vision Standard [AIA, Guide to understand vision standards]

Hardware: FPGA-based image processing products

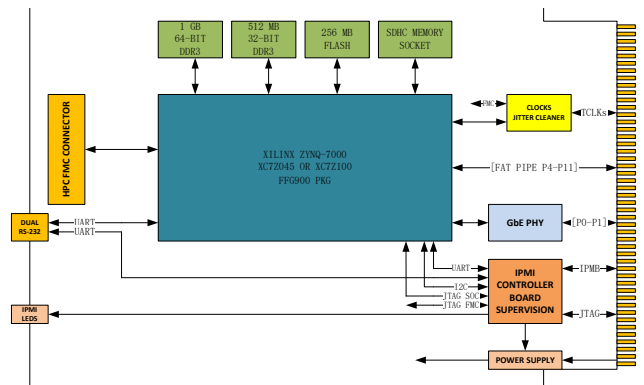
NAT's Vision platforms are built around a set of FPGA/CPU/GPU/DSP resource boards combined with a choice of interface boards and system platforms, providing the best fit solution for any application.

NAMC-ZYNQ-FMC

Single width, mid-size AMC (Advanced Mezzanine Card) with XILINX ZYNQ-7000 FPGA and an FMC slot per VITA57.1

The NAMC-ZYNQ-FMC carrier provides an onboard reconfigurable System-On-Chip (SoC) Xilinx Zynq-7000. The Zynq-7000 combines FPGA section with a dual core ARM Cortex A9 processor.

The SoC interfaces directly to the FMC slot via a High Pin Count Connector and to the MTCA back-plane offering PCIe, RapidIO (SRIO) or XAU (10 GbE).



Key features

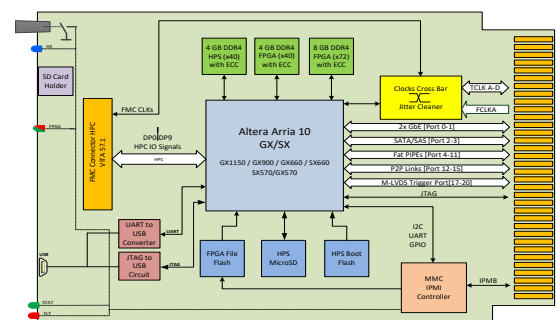
- Xilinx Zynq-7000 XC7Z045 or XC7Z100
- High-Pin Count FMC slot per VITA 57.1 compliant AMC.1, AMC.2 and AMC.4 compliant
- Dual banks of DDR3 memory accessible by FPGA & processor sections, allowing large buffer sizes and queuing during processing
- NAND and NOR flash memory for storage and application
- IPMI 2.0 compliant



NAMC-ARRIA10-FMC

Single width, full-size AMC (Advanced Mezzanine Card) with Altera Arria-10 GX or SX FPGA and a high pin count VITA57.1 compliant FMC slot.

The NAMC-ARRIA10-FMC carrier provides an on-board reconfigurable System-On-Chip (SoC) Altera Arria-10 FPGA, combining up to 1.5M logic elements, high speed SERDES interconnects and dual core ARM Cortex A9 cores.



Key features:

- Altera Arria 10: GX1150, GX900, GX660, GX570, SX660, SX570
- High-Pin Count FMC slot per VITA 57.1
- Up to 16GB of SDRAM
- Backplane Connectivity: GbE, SATA/SAS, PCIe, Serial RapidIO (SRIO), XAUI (10 GbE), Point-to-Point Links, Trigger lanes



FMC-4GigE-PoE

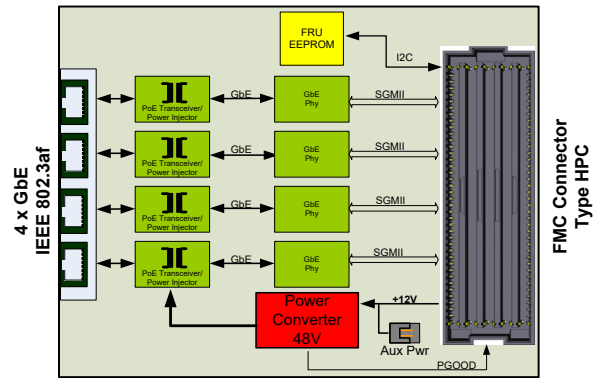
The FMC-4GigE-PoE module is an FMC Vita 57.1 standard compliant Ethernet switch module with support for Power over Ethernet (PoE).

It provides 4 front panel GbE ports compatible with IEEE802.3af. The board is capable to drive power to the Ethernet links per IEEE802.3af standard (max 15.4W per link, max 50W total)

Key features

- 4 IEEE 802.3af compatible front ports
- High efficient power converter
- 15.4W per PoE link
- Up to 50W total PoE power
- FMC HPC Connector

FMC-4GigE-PoE



NAMC-4GigE-PoE

The NAMC-4GigE-PoE is an MTCA compliant Ethernet switch module with support for Power over Ethernet (PoE).


It provides 4 front panel GbE connections compatible with IEEE802.3af (at). The board can drive power to 4 Ethernet links per IEEE802.3af standard (15.4W per link) or two links per IEEE802.3at standard (25.5W per link)

The four front panel Ethernet connections are aggregated to the MTCA backplane's 1GbE ports (0/1) or to the 10/40GbE fatpipe ports (4-7 or 8-11).

Key features

- 4 IEEE 802.3af compatible front ports
- High efficient power converter
- 15.4W per PoE link
- Up to 50W total PoE power
- Aggregation of 4 front panel GigE streams into one 10/40G backplane stream

AMC-4GigE-PoE

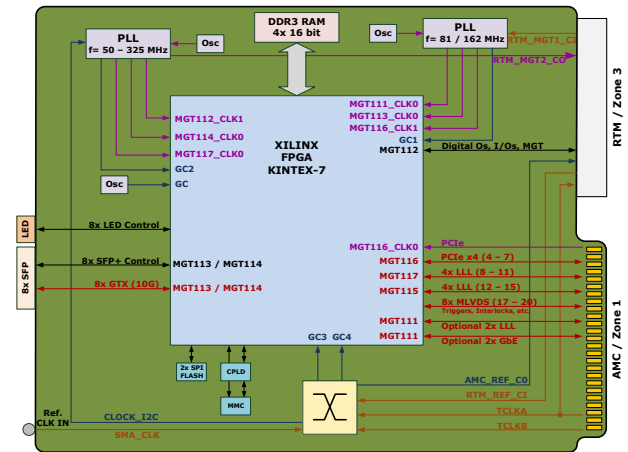


NAMC-TCK7

The NAMC-TCK7 is a high-performance, low-latency data processing AMC based on the Xilinx® Kintex™-7 FPGA, which provides industry-leading performance-per-watt.

The NAMC-TCK7 is suitable for a range of data processing applications from wireless network remote radio head to vision applications.

It provides a number of low latency links via the front panel, the backplane and a rear transition module (RTM), with up to 10 Gbps throughput each.



Key features

- Double-width, mid-size MicroTCA.4 AMC
- Xilinx® Kintex™-7 FPGA (XC7K355T or XC7K420T)
- 4x4Gb 1066 Mbps 64-bit DDR3 SDRAM
- 8x10 Gbps* SFP+ external interfaces
- Clock distribution circuit
- Supports 10 GigE Vision

* 12.5 Gbps possible with Xilinx® Kintex™-7 speed grade -3

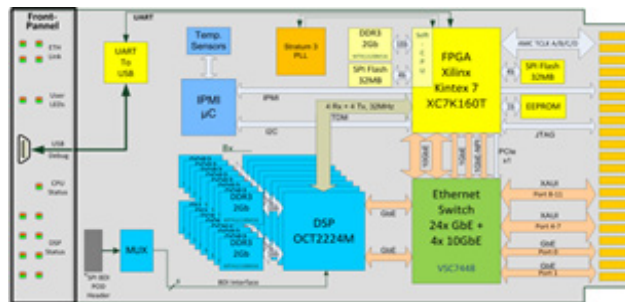


Hardware: Post-processing products

Audio and video acceleration card: NAMC-ODSP-M

The NAMC-ODSP-M adds video and audio acceleration to MicroTCA® (µTCA, MTCA) systems. It provides a broad range of audio/video codecs and algorithms, including H.264, MPEG4, H.263, and G.711, G.729, SILK and Opus, in a compact AMC module.

The integrated media gateway software provides rich audio and video applications with many codecs and protocols included and a simple configuration API so no DSP programming is required.



Key features

- Up to eight OCT2224M DSPs
- Xilinx Kintex-7 FPGA
- Redundant 1/10G Ethernet connectivity
- Integrated media gateway software
- Capable of up to:
 - 5,120 channels G.729AB/G.711 transcoding
 - 16 simultaneous H.264 HD/SD transcodes



Voice and video transcoding performance

	Channels per NAMC-ODSP-M1	Channels per NAMC-ODSP-M8
G.711 (20 ms) ↔ G.729AB (20 ms)	640	5.120
G.711 (20 ms) ↔ AMR NB (20 ms)	429	3.432
G.711 (20ms) ↔ Opus NB (20ms)	129	1.032
MPEG4/CIF → H.263/CIF 30 FPS	23	184
H.264/SD 30 FPS → H.264/720p30	2	16

NATvision: Post-processing on general-purpose CPUs

PowerPC	NAMC-QorIQ-xxxx: 8 different board configurations and up to eight Power-PC cores and data path acceleration architecture running embedded Linux
Intel x86 PC	NAT-MCH-COMex-E3: quad-core Intel® Xeon® E3 CPU running desktop Linux or Windows
	Third party AMC Intel CPUs
GPGPU	NVIDIA Tegra K1 AMC modules
	External PCs with NVIDIA graphic cards connected via optical PCIe or 100G connection to the MTCA system

Starter Kits

To speed the development of your machine vision applications, N.A.T. offers NATvision starter kits. Each starter kit contains everything required to start the development of a vision application:

- Chassis
- FPGA board with ARM CPU
- GigE-PoE front end
- Time-limited demo license of VisualApplets
- Pre-compiled demo application in source form

NATvision-mini

Our entry level system for vision applications, consisting of:

- NATIVE-Mini: 2-slot MTCA system with embedded MCH
- FPGA resource board: NAMC-ZYNQ-FMC FPGA board
- FMC-4GigE-PoE: 4 times GigE-PoE Interface board
- GigE Vision FPGA IP core
- Linux BSP running on ARM Core
- Vision demo application ready-to-go



NATvision-mini+

NATvision-mini platform enhanced by a standard CPU board

- NATIVE-Mini: 2-slot MTCA system with embedded MCH
- FPGA resource board: NAMC-ZYNQ-FMC FPGA board
- FMC-4GigE-PoE: 4 times GigE-PoE Interface board
- GigE Vision FPGA IP core
- Linux BSP running on ARM Core
- Vision Demo application ready to go
- Intel® multicore CPU board running Windows and desktop Linux





About N.A.T.

Founded in 1990 with the aim of developing high-performance network interfaces for industrial computers, N.A.T. is a privately owned and financed company with headquarters in Bonn, Germany and certified distributors and sales agents worldwide.

The company has developed substantial knowledge in networking technologies across a wide range of open, standards-based architectures including VMEbus, CompactPCI, PCI, PCIe, PMC, Industry Pack Modules, and M-Modules. N.A.T. was at the forefront when the AdvancedTCA, MicroTCA and AMC standards were introduced and is today one of the leading suppliers for board and system level products based on AMC and MicroTCA.

The product portfolio includes line interfaces, network processors, single- and multicore data engines, management and system controllers, 19" rack mountable chassis, power supplies and of course the communication protocols and middle-ware to build turn-key and application-ready systems.

For more information, please visit www.nateurope.com.

N.A.T. GmbH | Konrad-Zuse-Platz 9 | 53227 Bonn | Germany |

Fon: +49 228 965 864 0 | Fax: +49 228 965 864 10

info@nateurope.com | www.nateurope.com