

Image Processing at Gigapixels per Second

White Paper

Executive Summary

Advanced vision systems such as 3D imaging, virtual reality, and broadcasting rely on one or multiple high-resolution, high-speed cameras. These camera systems capture and deliver multiple gigapixels per second that need to be processed and/or compressed in real-time.

While camera sensor technology and supporting interfaces are quickly evolving, traditional PCIe/CPU/GPU-based PC-architectures are lacking the required performance to capture, process, and store images at such data rates.

Frame grabbers based on high-end FPGAs with ultra-fast transceivers and massive on-board memory bandwidth offer the necessary infrastructure for real-time processing and compression. Open-FPGA architectures allow developers to customize their acquisition path and embed their own image processing algorithms and compression blocks.

With supporting eco-system and development tools like Gidel's ProcVision IDE, the FPGA code development is greatly simplified and accelerated and no longer necessitate deep FPGA expertise. Thus, vision system designers can now implement custom cost-effective gigapixel-solutions with unprecedented performance.

The processing of video streams at 10 to 100 Gbit/s and beyond benefits from novel processing hardware architectures.

Image Sensors are Setting the Pace

Advances in CMOS image sensor technologies have enabled multi-mega pixel imagers with frame rates of hundreds to thousands of FPS at cost-effective prices. Companies like Gpixel, Luxima Technology, Teledyne e2v, AMS/CMOSIS, ON Semiconductor, and Sony are making significant contributions to this development (see figure 1).

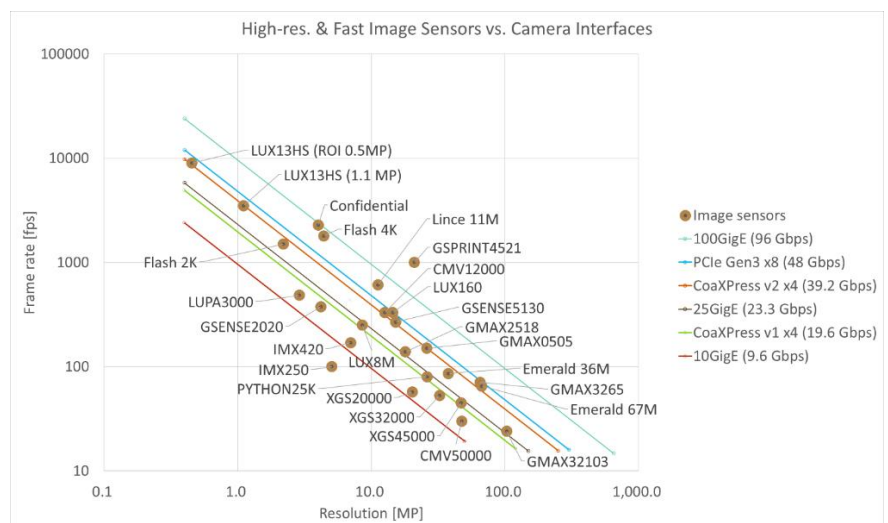


Figure 1: Resolution and frame rates of high-end image sensors over the effective bandwidth of common video interface standards (Source: Vision Markets GmbH)

The next generation of image sensors will generate data rates of 160 gigabits per second (Gbps) and beyond. Furthermore, multi-camera applications have become ubiquitous, especially in areas from Virtual Reality to Broadcasting, Surveillance, Medical Imaging, and Quality Inspection in 3D or with high resolution. For example, a 3D sports broadcasting system, may comprise more than 30 cameras of 65 MP resolution each at 30 fps.

Both, latest high-end image sensors and multi-camera applications deliver multi-gigapixels per seconds and several 100 Gbps respectively. These enormous data rates need to be captured, pre-processed, analyzed, and often also compressed and stored in real-time with high-precision synchronization and low latency – a requirement that exceeds the capabilities of CPUs based architectures by far. Instead, such demand can only be met by novel heterogenous processing solutions utilizing the unique capabilities of FPGAs, GPUs and/or CPUs.

High Bandwidth Challenges

For the transmission of sensor data rates beyond 20 Gbps, only few options among standard camera interfaces exist: 25, 50, or 100GigE, multi-link CoaX-Press v2, and PCIe. At 20+ Gbps, fiber optic cables replace copper cables to extend the transmission distance from 25 m to up to 40 km.

An additional challenge of high-bandwidth imaging lies in the transmission of video stream to the high-performance processor of the host, be it a GPU, FPGA and/or CPU. The motherboard interface of video capture cards is typically PCIe Gen. 3 x8 with an effective bandwidth of merely 48 Gbps. Moreover, within the host processing system, the CPU/GPU and the RAM bridge between the graphics card and the main memory must operate sufficiently fast to avoid frame loss. Smart NICs succeed in distributing peak load on the PCIe and significantly reducing the workload on the host CPU, yet it often comes at the expense of lost image frames due to insufficient processing power.

Solutions for the Real-time Processing of Gigapixels per Second

At data rates of tens and hundreds of Gbps, except for expensive ASICs, only frame grabber architectures based on high-end FPGAs provide the necessary processing performance to overcome the mentioned challenges. These grabber cards need to go far beyond the traditional pre-processing steps to perform complex imaging algorithms from wavelet transformations all the way to deep learning inference and real-time compression. Compression is a mandatory feature to overcome the PCIe and host memory bandwidth bottlenecks.

The design of such high-end frame grabber is a challenge by its own particularly when it comes to the implementation of algorithms utilizing data from several image regions or multiple sensors. To circumvent possible bottlenecks and to enable flexibility for distributed processing, the frame grabber must include powerful transceivers, sufficient FPGA resources, high on-board memory access, and fast DMA offload engines. Such frame grabbers typically

include FPGA internal memory with access rates of TB/s, and 10+ GB of DDR4 on-board memory with access rates of hundreds of GB/s.

The implementation of machine vision algorithms on FPGAs generally requires in-depth expertise in FPGA programming. Moreover, implementing multi-camera acquisition and processing on a single FPGA requires the integration of multi-interfaces, camera protocol, multi-source processing algorithms, memory controller, I/O port control, and a host bridge. Besides a performance-optimized architecture, it is critical that the frame grabber will be supported by an Integrated Development Environment (IDE) that enables non-FPGA-experts to develop the imaging algorithms and to integrate multiple FPGA functional blocks.

Open-FPGA Frame Grabbers Optimized for High-Bandwidth

With nearly three decades of experience, Gidel, an Israeli technology leader, has created an eco-system of off-the-shelf frame grabbers optimized for ultra-high bandwidth and multi-sensor acquisition that allows the developers to add their proprietary algorithm code to the existing grabbing pipeline. Thanks to a dedicated development suite, adding image processing algorithms and customizing the acquisition path is simple and can be performed even by non-FPGA experts. Gidel's development suite significantly accelerates system development without compromising performance.

Gidel's PCIe frame grabbers, modules, and carrier boards allow vision system designers to leverage the latest advancements in FPGA technology, such as Intel's Stratix 10 and Arria 10 series. Gidel's latest Proc10N module can grab and process up to 4 x 100 GigE cameras or 16 x 10 GigE cameras simultaneously with accurate low-latency synchronization. With 300 GB/s access to DRAM, the Proc10N enables real-time processing even for the most bandwidth-demanding applications. The Stratix 10 NX features exceptional matrix computation capabilities with dedicated Tensor-Blocks ideal for high-performance inference computation, including complex Deep Learning Networks.



North America:

1600 Wyatt Drive Suite 1
Santa Clara, CA 95054
+1-408-969-0389
sales_usa@gidel.com

International:

2 Ha'ilan St., Northern Ind. Zone
POB 281, Or Akiva, Israel 3060000
+972-4-610-2500
sales_eu@gidel.com

www.gidel.com