High-performance image processing using FPGAs

Successful technology transfer of GigE and implementation of 10 GigE Vision protocol

High-performance video data acquisition, on-the-fly image processing and real-time data evaluation have gained tremendous importance in areas such as industrial automation, process monitoring or quality inspection. For accelerator control or in photon science experiments, rapid 2D image systems are crucial assets used for fast feedbacks, target protection or *in-situ* parameterisation of experiments. Data analysis meanwhile is carried out by artificial-intelligence algorithms requiring large data throughput combined with sufficient parallel computing power, e.g. using field programmable gate arrays (FPGAs). In this article, we present our implementation of the GigE Vision protocol in FPGA, offering precise control of the execution flow. The implementation is able to acquire data from commercial off-the-shelf cameras from various vendors and allows high-speed image capture, real-time processing and distribution of the information to other components in the system.

Introduction

GigE Vision is a communication protocol for transferring images and video over Gigabit and 10 Gigabit Ethernet. It is one of the most widely used protocols for connecting cameras to frame grabbers and other receivers.

Because of their internal structure, FPGAs are well suited for demanding input/output tasks and provide precise control of the execution flow compared to CPUs and GPUs. The implementation of the GigE Vision protocol in FPGA allows very high-throughput, low-latency reception of images and tight integration with image-processing algorithms.

GigE Vision specification

Strict adherence to the GigE Vision standard improves interoperability as well as user experience when connecting cameras from different vendors. To confirm the compliance of our implementation with the GigE Vision standard, we attended a plugfest hosted by the Automated Vision Association (AIA) in 2019. Our setup successfully interoperated with cameras from Allied Vision, Basler, JAI, Hamamatsu and Teledyne, confirming the compliance and interoperability with other vendor products. Meanwhile, the GigE Vision certification has been granted.

Technology stack

The right side of Fig. 1 shows the technology software and firmware stack from reception to processing of video frames. For clarity, some details are omitted. Displayed at the bottom is a GigE or 10 GigE Vision-compliant camera equipped with an Ethernet port. The camera is directly connected to an FPGA, where the Ethernet physical layer is handled by vendor-provided intellectual property (IP) cores. The data stream is then forwarded to a network stack, consisting of an Ethernet MAC and a UDP/IPv4 engine. In the next step, the data is transferred to the GigE Vision core, which extracts the image from the GigE Vision data packets. Afterwards, the image data is provided to the application over a standardised interface, the AXI4 Stream Video. Using a standardised interface increases modularity and simplifies interfacing to different modules, e.g. video direct memory access (DMA), custom logic and modules written in Vivado High-Level Synthesis (HLS) or Matlab. The data is then transferred to the CPU, either via PCI Express or internal interconnects in Xilinx Zynq devices. On the CPU, a server provides a high-level interface for the firmware. A Python library handles the application part, together with a GUI to display the images on the screen.

Performance test

To illustrate the performance, we carried out a test with a FLIR Oryx 10 GigE Vision camera (model ORX-10G-51S5C) together with our GigE Vision stack running on powerful FPGA module with a Kintex 7 FPGA (DAMC-TCK7). We configured the frame size to 1920 x 2048 pixels in RGB8 format (24-bit per pixels). The exposure time was reduced to reach a frame rate of up to 82.38 frames per second measured with an AXI4 Performance Monitor, which corresponds to a data flow of 8288 Mb/s to a DDR3 memory. Most electronic boards are equipped with fast DDR3 or DDR4 memory, the fastest being Struck SIS8160 with a total throughput of 205.3 Gb/s [1]. If hardware permits, frame rates of 1 kHz and higher can be achieved, which is of particular interest for the new DESY high-power laser system KALDERA.

Applications

Supported hardware electronic boards

The GigE Vision stack can be implemented on any Xilinx FPGA of decent size with sufficiently wide memory

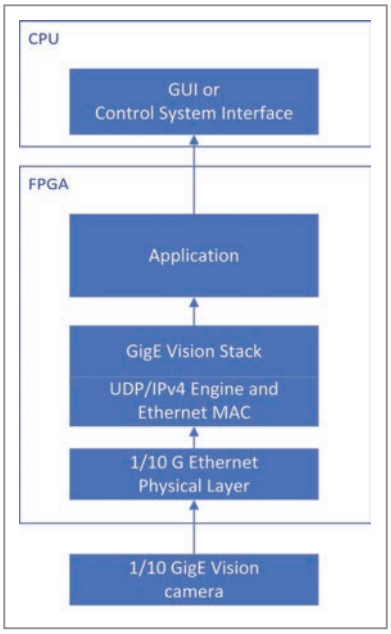

Figure 1

Left: GigE Vision application in MicroTCA crate. Right: Technology stack

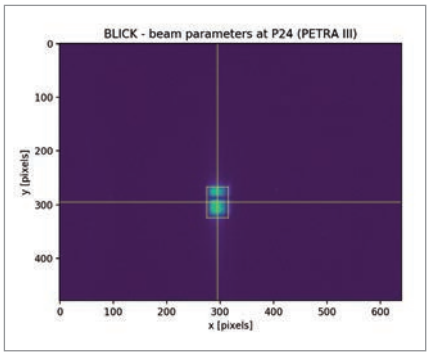

Figure 2

Output from the BLICK software

connection. So far, the stack was successfully ported and tested on the following electronic boards:

- DAMC-TCK7 licensed to N.A.T
- DAMC-FMC2ZUP licensed to CAENels
- NAMC-ZYNQ-FMC N.A.T
- SIS8160 Struck GmbH
- ZCU102 Evaluation Kit Xilinx
- KCU116 Evaluation Kit Xilinx

Several of the supported boards are advanced mezzanine cards (AMCs) compatible with the MicroTCA platform. For image acquisition and processing, the MicroTCA platform provides superior characteristics, such as:

- Aggregation of a large number of cameras [2]
- Information from the cameras can be combined with other sensors and used for fast feedback
- CPU with application software can be integrated in the crate, using PCI Express on the backplane
- Timing on the backplane to timestamp the frames
- Future developments will also allow the camera acquisition to be triggered from the backplane (e.g. from a global timing system or from other cards in the crate)

BLICK – BeamLine Instrumentation Camera Kit

BLICK is an application and a software tool kit that uses the GigE Vision stack to provide simultaneous readout of eight cameras with real-time diagnostic capabilities (e.g. 2D mean and standard deviation [3]), allowing the size and position of beams (electron, laser or X-rays) to be determined.

A Python-based GUI provides a visualisation of the images, and results of the real-time diagnostic module are displayed. An easy interface allows users to program specific functions themselves (Fig. 2). In future, BLICK will also be integrated in the areaDetector [4] framework, providing data to the Experimental Physics and Industrial Control System (EPICS).

Use cases

The GigE Vision stack can be used for any 2D imaging application where fast, deterministic or real-time data processing is needed. The stack is used as part of the European Spallation Source (ESS) target protection system [5] and at the PETRA III experiments (beamline P24) at DESY. It is of particular interest when frame rates of 100 Hz to 1 kHz are required or low latency is mandatory, e.g. for fast feedbacks.

The development has a high technology transfer potential. The MicroTCA Technology Lab at DESY has licensed the GigE Vision IP core to several industry customers and provides customisation regarding online processing algorithms, implementation and porting to new hardware platforms. Algorithms are written in C++ and compiled to FPGA logic with Vivado HLS, allowing for easy and fast developments.

Contact: Jan Marjanovic, jan.marjanovic@desy.de
Sven Stubbe, sven.stubbe@desy.de
https://techlab.desy.de/products/gige_vision_solutions/

References:

[1] J. Marjanovic, S. Stubbe, A. Gornott, “Update on GigE Vision implementation in FPGA”, ARD-ST3 2019, https://indico.desy.de/indico/event/23131/session/6/contribution/46
[2] V. Dirksen, “Up to 48 GigE-Vision Cameras benefitting from the MicroTCA architecture”, MTCAWS2018, https://indico.desy.de/indico/event/20703/session/14/contribution/56
[3] J. Marjanovic, “Low vs High Level Programming for FPGA”, IBIC 2018, http://ibic2018.vrws.de/papers/thoa01.pdf
[4] areaDetector, https://areadetector.github.io
[5] M. Donna et al., “Optical System for ESS Target Protection”, IBIC 2014, http://accelconf.web.cern.ch/AccelConf/IBIC2014/papers/tupf27.pdf