Overview

Full GPU Image Processing Pipeline for Camera Applications

Serzhenko Fyodor, Uvarov Andrey

Fastvideo, Russia

Problem to Solve

Fastvideo

Real time image processing

- Industrial cameras with high resolution and high frame rate
- Multiple camera systems
- Broadcasting solutions
- High speed cameras

What is common for all cases: data rate up to ~GPix/s

Image Processing in Realtime

Fastvideo

Modern machine vision cameras

- Image sensor resolution from VGA to 30 Mpix
- Frame rate up to 500 fps
- Bit depth: 8 – 12 bits
- High speed interface to PC
- Multiple camera systems

Key Points

Fastvideo

- Implementation of full image processing pipeline on GPU for real time camera applications
- We can do that faster and with better quality in comparison with solutions on CPU, DSP, FPGA
- Image Processing SDK from Fastvideo

Possible Solutions for Image Processing

Fastvideo

- In-Camera HW processing on FPGA, DSP, CPU
- Image processing on frame grabber
- Image processing on CPU with SSE/AVX
- Image processing on GPU

Basic Image Processing Pipeline

Fastvideo

- Image Acquisition
- Preprocessing
- Demosaicing
- Resize / Sharp / OpenGL output
- JPEG compression

Preliminary Stages

Fastvideo

- In-Camera HW image processing
- Camera-to-PC data transfer over cables
- Image acquisition on PC side
  - Data transfer to CPU memory
  - RDMA GPUDirect

Image Preprocessing on GPU

Fastvideo

- Image data unpacking and rearranging
- Fixed-Pattern Noise subtraction
- Vignetting
- Bad pixels removal
- White balance

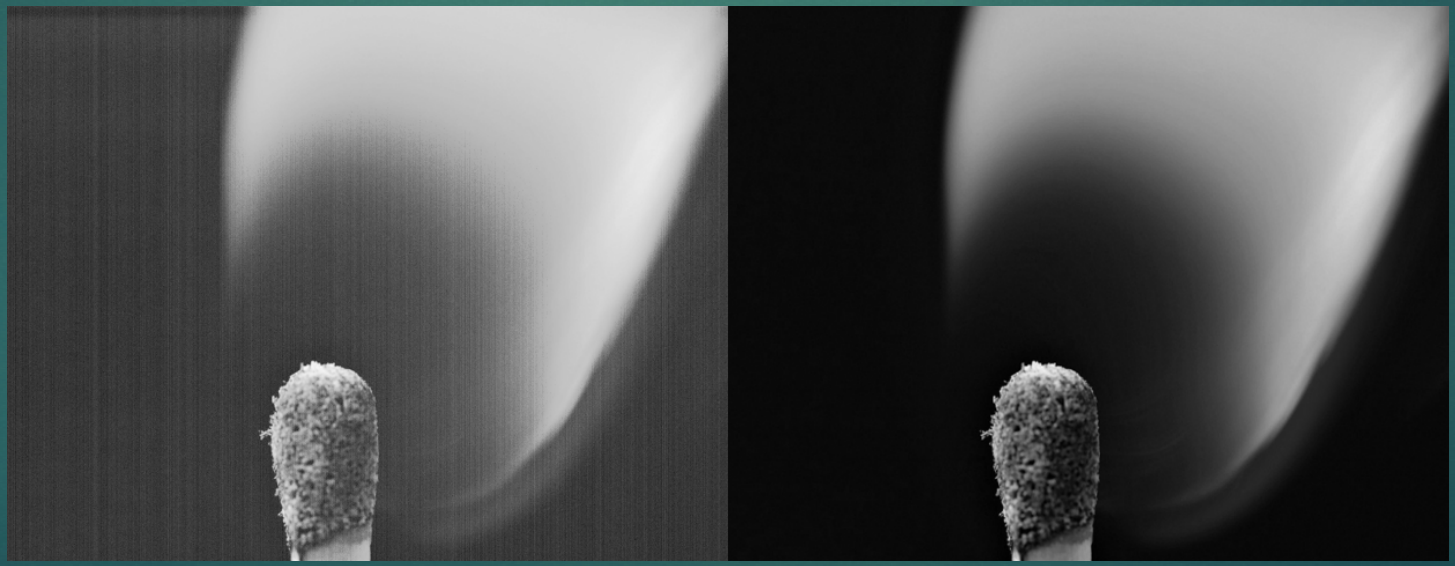

Image Demosaicing (Debayer)

Fastvideo

Demosaicing means image color interpolation from Bayer image sensor

(input data 8 bits, output data 24 bits)

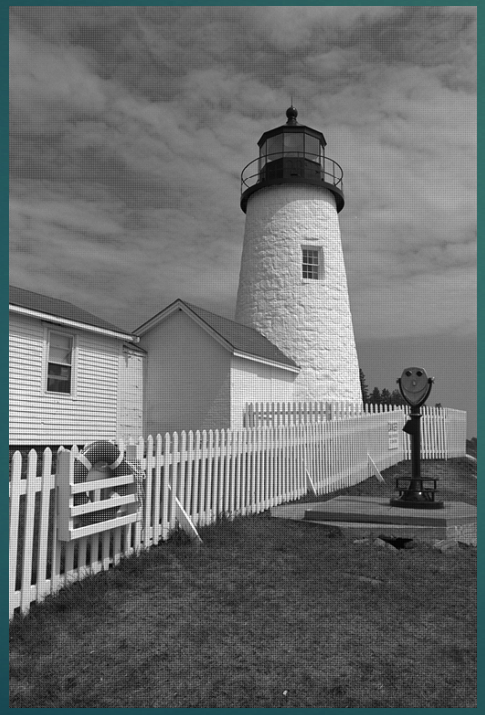

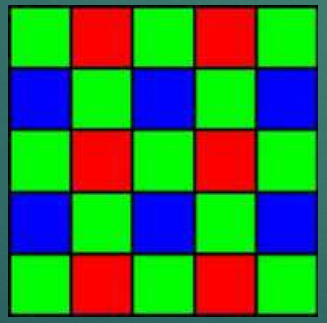

Performance of Demosaic Algorithms

Fastvideo

- Bilinear – bilinear interpolation for each pixel
- HQLI – High Quality Linear Interpolation with 5x5 kernels
- AHD – Adaptive-Homogeneity-Directed algorithm
- DFPD – Directional Filtering with a Posteriori Decision

	Bilinear	HQLI	AHD	DFPD
Fastvideo on Titan GPU		2400 MB/s		1770 MB/s
IPP on CPU Core i7 3770	2800 MB/s		14 MB/s	
GPU kernel performance		21000 MB/s		6000 MB/s
SSIM	0.873	0.965	0.968	0.978
PSNR (dB)	30.4	36	37.4	39

CUDA JPEG Codec

Fastvideo

In 2011 we have created the first fully parallel CUDA JPEG codec.

Main idea: JPEG algorithm consists of consecutive stages and it's applied to blocks 8x8 or 16x16 pixels. Parallel block handling for all stages is a must for fast JPEG codec.

Key for fast decoding: presence of Restart Markers in bitstream.

Results for GeForce GTX Titan

JPEG Encoding: 2500 – 5500 MB/s, JPEG Decoding: 2000 – 5000 MB/s

CUDA JPEG results for 580/680/Titan

Fastvideo

- Host-to-Device and Device-to-Host transfers are included
- Win-7 (32-bit), NVIDIA drivers 332.21

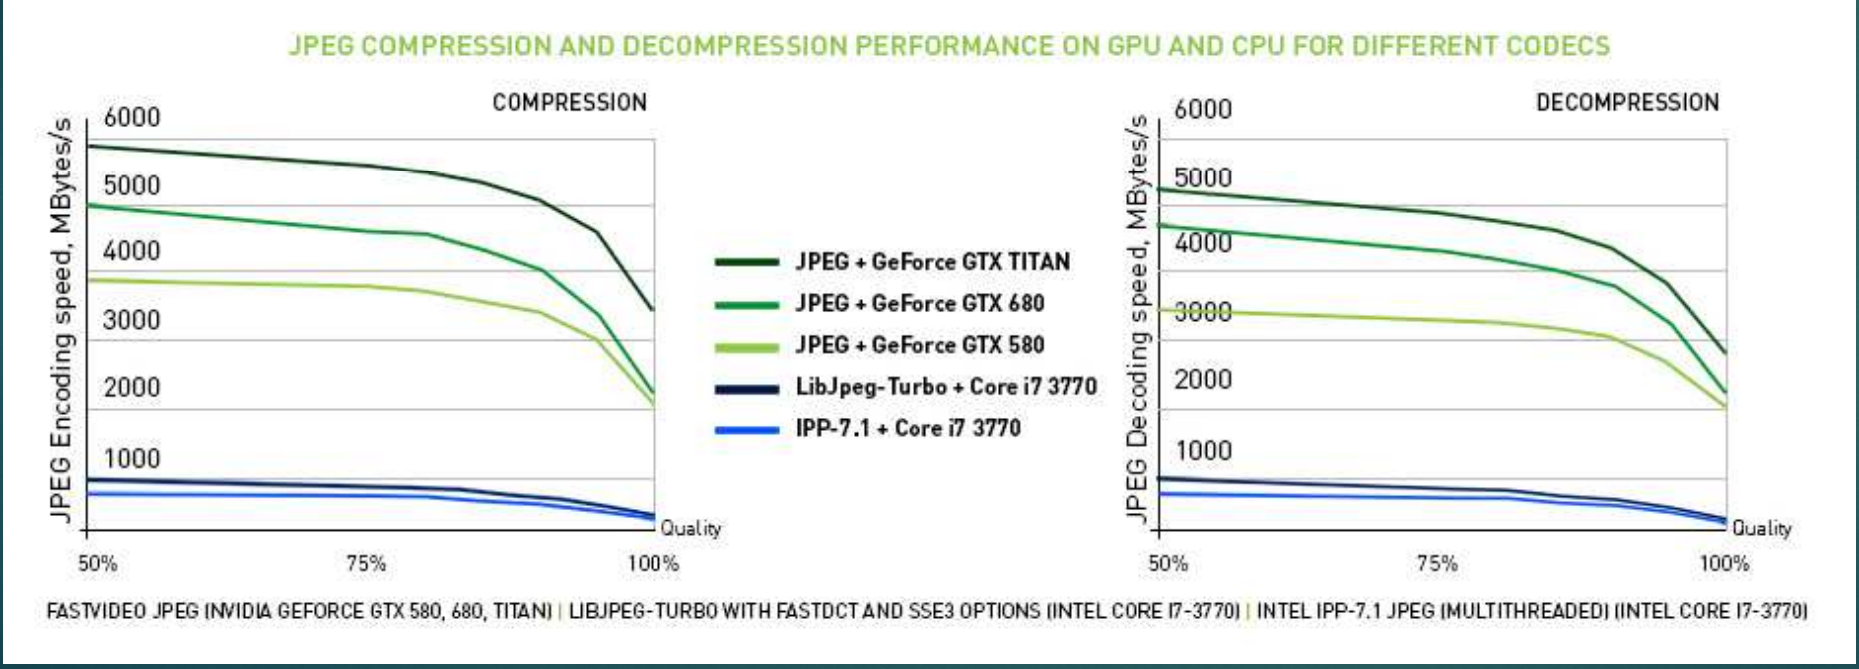

JPEG and Image Distortion

Fastvideo

- JPEG is not always guilty for insufficient image quality
- There are other lossy stages in the pipeline:
  - Preprocessing
  - Demosaicing
  - Denoising
  - LUT
  - Resize / Sharp

Solution: use “visually lossless” JPEG quality 85-100% and always check quality losses for other stages of the pipeline

Fastvideo

Combined Demosaic + JPEG

- Data after demosaic are tripled
- Transfer time over PCI-Express bus is important
- Combined solution saves time for CPU-GPU transfers
- Low-quality demosaic is unacceptable

Benchmarks for Demosaic and JPEG Encoder

	1920 x 1080	2048 x 2048	4096 x 4096
Demosaic (DFPD)	1.6 ms	3.0 ms	10.2 ms
JPEG (90%, 4:2:0)	2.3 ms	3.2 ms	10.5 ms
Demosaic + JPEG	2.6 ms	3.7 ms	11.7 ms

Final Benchmark on GPU (Titan)

Fastvideo

- CMOSIS image sensor CMV20000, 5120x3840, 12-bit, 30 fps
- GeForce GTX Titan GPU
- Host to device transfer ~1.5 ms
- Demosaic ~3.1 ms
- JPEG encoding (90%, 4:2:0) ~7.8 ms
- Device to Host transfer ~1.3 ms
- Total: ~13.7 ms

P.S. This is the benchmark for PCIE camera CB-200 from XIMEA

Camera Applications

Fastvideo

- Machine vision and computer vision
- XIMEA solutions – www.ximea.com
- Imperx solutions – www.imperx.com
- Multi-camera image processing on single GPU
- Real time image processing for broadcasting and surveillance
- Video conferencing
- High speed imaging

Non-Camera Applications

Fastvideo

- Image processing speedup to image viewers, photo editors
- High performance batch image processing
- Fast resize and thumbnail generation
- Cloud solutions for image processing
- Medical imaging

Solution for Photo Hosting

Fastvideo

Task description: **load-decode-resize-encode-store**

- Image load ~1.5 ms for 2048x2048 jpg image
- JPEG decoding ~3.4 ms
- Downsize to 1024x1024 with bicubic algorithm ~0.7 ms
- JPEG encoding (quality 90%, 4:4:4) ~3.4 ms
- Image store ~1.0 ms
- GPU processing time ~7.5 ms
- Total time ~10 ms

What we could offer to our customers

Fastvideo

- Standard SDK
  - Preprocessing functions
  - Demosaic
  - JPEG Encoder and Decoder
- Custom SDK for particular task
- Algorithm design
- Software optimization

More Features of SDK

Fastvideo

- Image acquisition module from a camera
- GPU Direct option to get images on GPU from a camera
- Unpacking module for various raw formats
- Wavelet denoising
- Color correction
- Image tiling and batch processing
- Crop/Resize/Sharp and pyramid images
- OpenGL output
- MJPEG integration with FFMPEG

Long Term Strategy

Fastvideo

- More image processing algorithms in our SDK
- More complicated parallel algorithms with higher quality
- JPEG2000 and more lightweight custom codecs
- Code optimization, use of the latest GPUs and APIs
- Parallel algorithm design and implementation

Links & Contacts

Fastvideo

- Fastvideo site for GPU image processing - www.fastcompression.com
- Contacts: info@fastcompression.com

Camera manufacturers with GPU image processing software

- XIMEA GmbH – www.ximea.com
- Imperx Inc. – www.imperx.com