NVIDIA OPTICAL FLOW

Abhijit Patait, 3/18/2019

NVIDIA

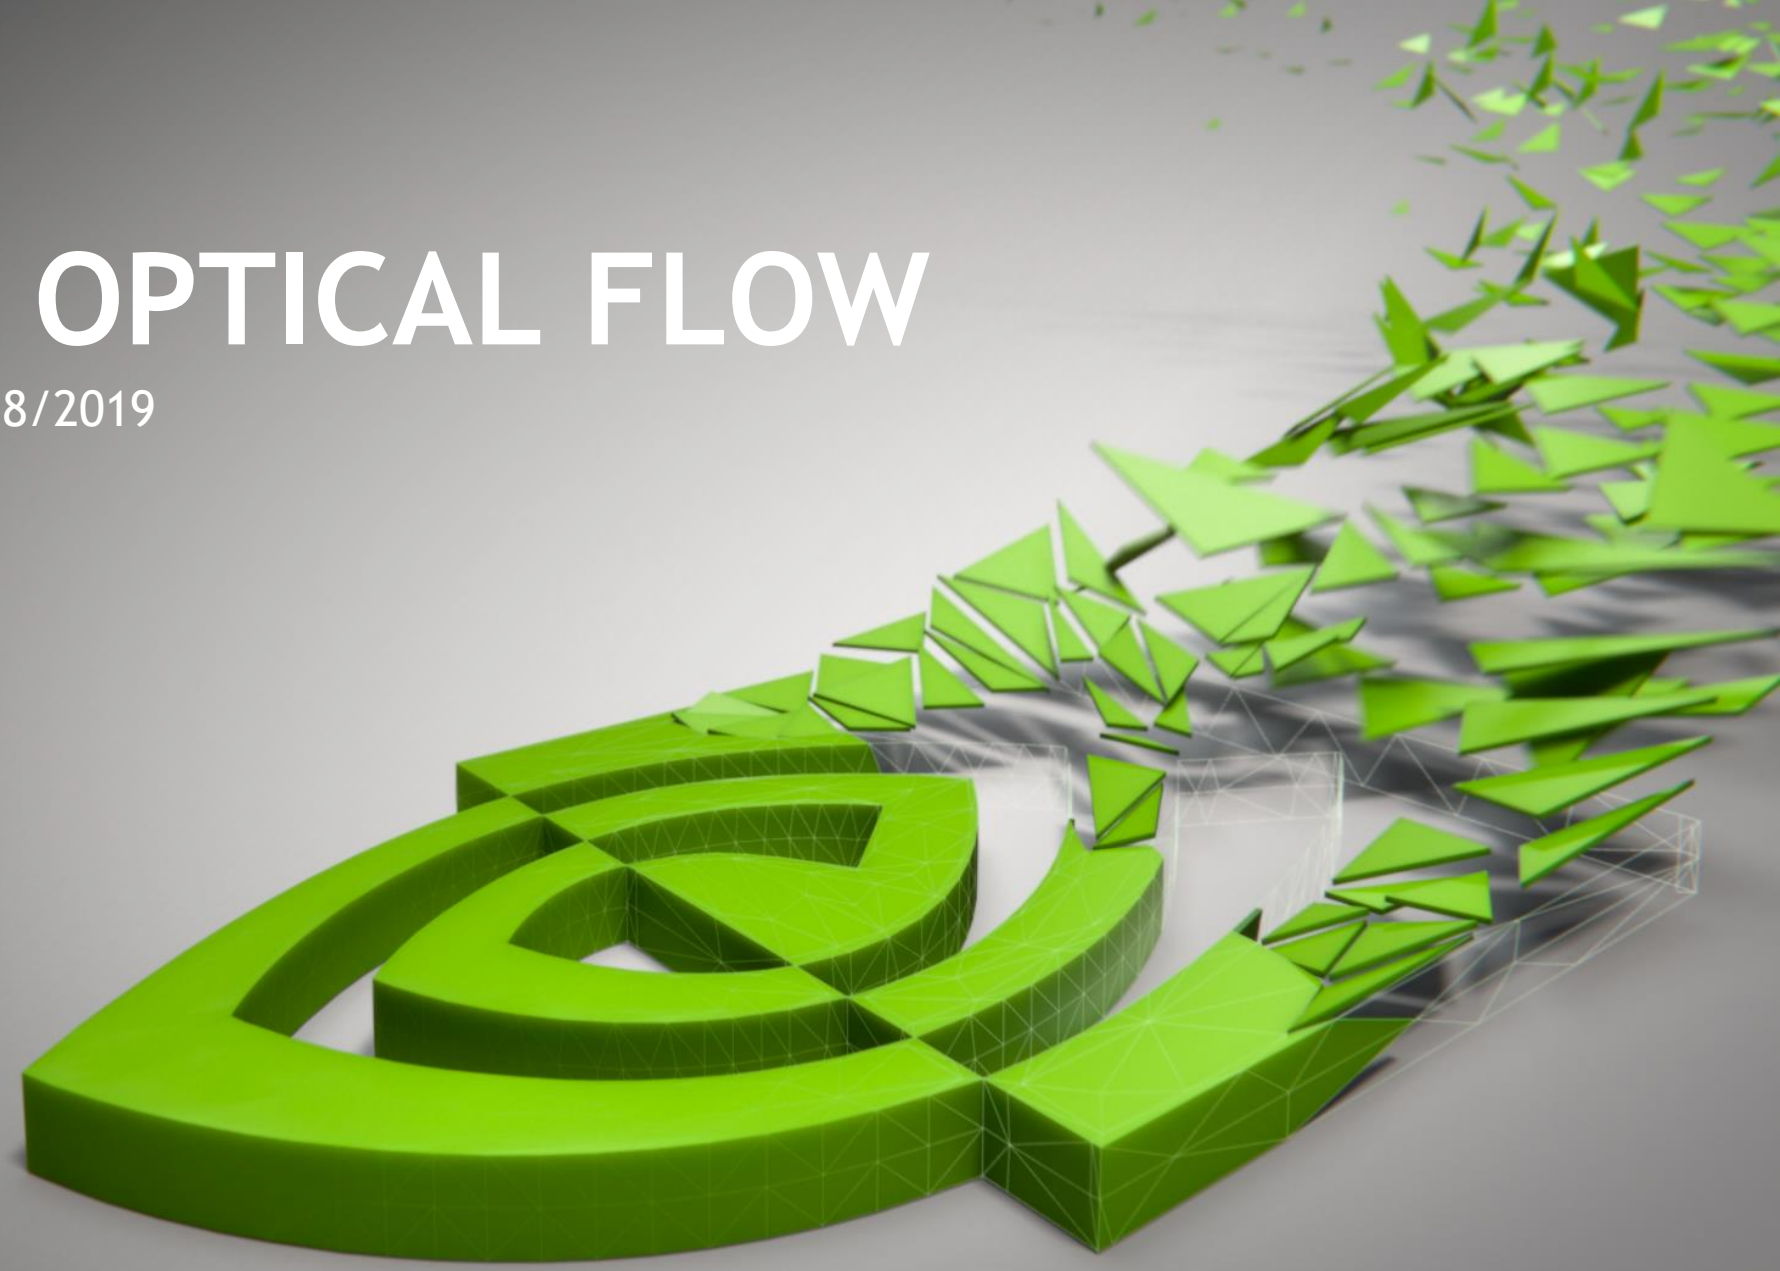

AGENDA

Optical Flow in Turing GPUs

NVIDIA Optical Flow SDK

Benchmarks

End-to-end applications

Roadmap

BACKGROUND


CAPTURE
C
CAPTURE
A
49
109
GavinBiscuit: Destroyed Counter-UAV
zukowski: called in UAV
zukowski: Captured Zone C
[heck]RyuukiAz
[USA]IRONMAN
Chaotic
zukowski
150
X
E
550
60
CORDITE


CAPTURE
C
CAPTURE
A
49
109
[heck]RyuukiAz
[USAJIRONMAN
Chaotic
zukowski
150
X
E
60
CORDITE
550

ESTIMATING PIXEL MOTION

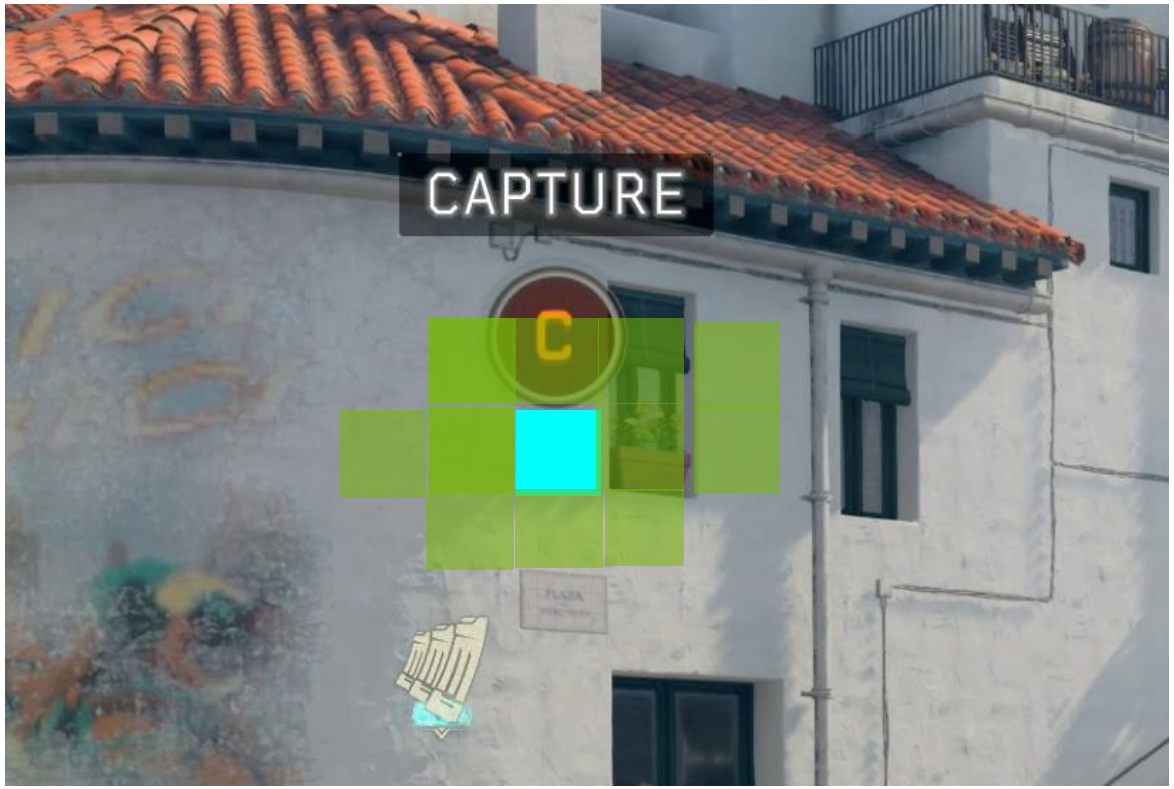

- "Video" motion vectors
  - Minimize encoding cost
  - SAD, SATD, RDO, intra modes, partitions
- Optical flow vectors
  - Visual motion
  - Current and surrounding pixels/blocks

ESTIMATING PIXEL MOTION USING NV GPUS

- ME-only mode – Maxwell, Pascal, Volta
  - Optimized for encoding – up to 8×8 granularity **motion vectors**
  - Video Codec SDK 7.0+
- Optical flow (OF) – Turing & beyond
  - New hardware in NVENC
  - Optical flow and stereo disparity
  - Optical Flow SDK 1.0 (released Feb 2019)

OPTICAL FLOW ENGINE

Capabilities

- Hardware
  - Up to 150* fps at 4K
  - 4 × 4 pixel granularity
  - ¼ pixel resolution
  - Accuracy comparable to best DL methods
  - Advanced algorithms to find true flow vectors
- Software
  - SDK (Windows, Linux, CUDA, DirectX)

**Dependent on device clock speed*

INTENSITY DIFFERENCES

Optical flow must be insensitive to intensity

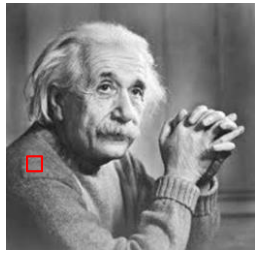

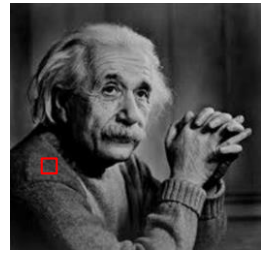

136	118	26	31	39
110	115	33	40	30
98	102	78	67	45
48	57	23	221	112
39	86	99	155	200

70	62	14	16	20
58	59	17	20	15
49	56	40	33	23
24	29	12	112	62
20	43	55	78	111

TURING OPTICAL FLOW VS MOTION VECTORS

	Turing Optical Flow	Pascal/Volta Motion Vectors
Granularity	Up to 4x4	Up to 8x8
Algorithm used	Visual motion optimization	Encoding cost optimization
Quality	Robust to intensity changes	Sensitive to intensity changes
Accuracy	Close to true motion Low average EPE (end-point error)	May deviate from true motion Higher EPE

NVIDIA OPTICAL FLOW SDK

NVIDIA OPTICAL FLOW SDK

- New Optical Flow C-API
- Scalable, accommodates needs of future hardware
- Linux, Windows 8.1, 10, server, …
- DirectX, CUDA interoperability
- OpenCV
- Public released – Feb 2019
- Legacy ME-only mode API continues to be supported

OPTICAL FLOW API

Basic functionality

Main Functionality (nvOpticalFlowCommon.h)

```
NV_OF_STATUS(NVOFAPI* PFNNVOFINIT) (NvOFHandle hOf, const
NV_OF_INIT_PARAMS *initParams);

NV_OF_STATUS(NVOFAPI* PFNNVOFEXECUTE) (NvOFHandle hOf, const
NV_OF_EXECUTE_INPUT_PARAMS *executeInParams,
NV_OF_EXECUTE_OUTPUT_PARAMS *executeOutParams);

typedef NV_OF_STATUS(NVOFAPI* PFNNVOFDESTROY) (NvOFHandle hOf);
```

CUDA and DirectX buffer management

```
nvOpticalFlowCuda.h & nvOpticalFlowD3D11.h
```

REUSABLE CLASSES

NvOF

Base class for all core functionality

NvOFCUDA

Input and output in CUDA buffers

NvOFD3D11

Input and output in DirectX buffers

USE VIA OPENCV

```
Mat frameL = imread(pathL, IMREAD_GRAYSCALE);

Mat frameR = imread(pathR, IMREAD_GRAYSCALE);

GpuMat d_flowL(frameL), d_flowR(frameR), d_flow;

Mat flowx, flowy, flowxy;

int gpuId = 0;

int width = frameL.size().width, height = frameL.size().height;


Ptr<cuda::NvidiaOpticalFlow_2_0> OptFlow = cuda::NvidiaOpticalFlow_2_0::create(perfPreset,
NVIDIA_OPTICAL_FLOW_HINT_DEFAULT, gpuId);
OptFlow->calc(d_flowL, d_flowR, d_flow);
d_flow.download(flowxy);
```

BENCHMARKS

OPTICAL FLOW QUALITY

Evaluation Methodology

- **Objective quality**
  - KITTI 2012/2015, Sintel, Middlebury
  - Average end point error (EPE)
  - Percentage of outliers – background, foreground and all
- **Subjective quality**
  - Flow maps
  - Frame-rate-up-conversion (video interpolation)

OPTICAL FLOW QUALITY

EPE - KITTI 2015

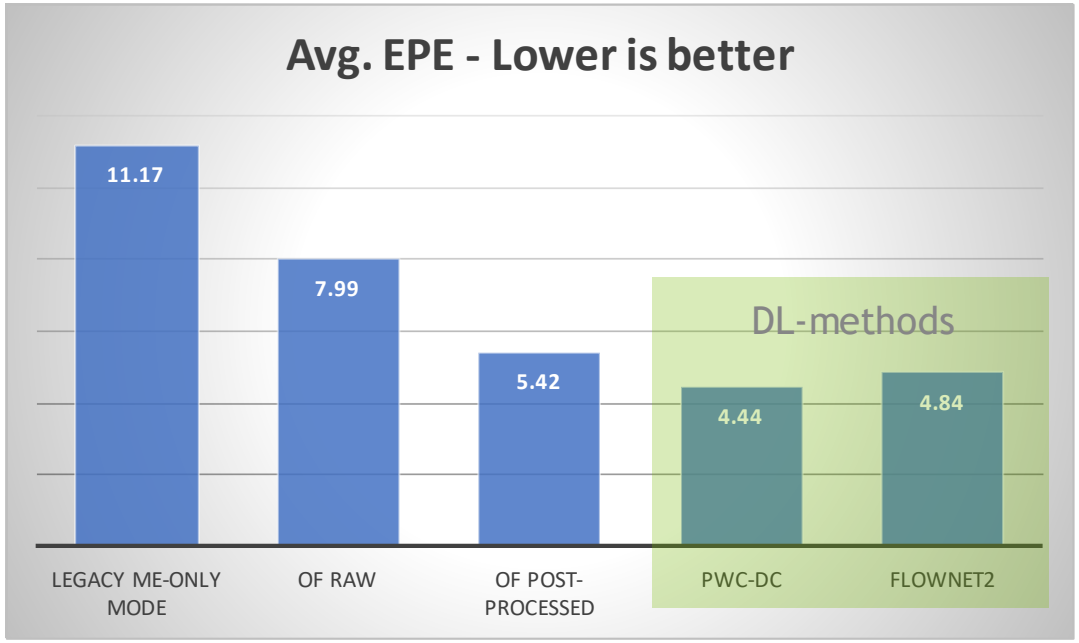

- EPE = End-point error = Euclidian distance between OF vector & ground truth
- Non-occluded EPE
- Occluded EPE higher but same trend
- KITTI 2012 EPE = 2.31
- Sintel EPE = 8

OPTICAL FLOW QUALITY

Outliers - KITTI 2015

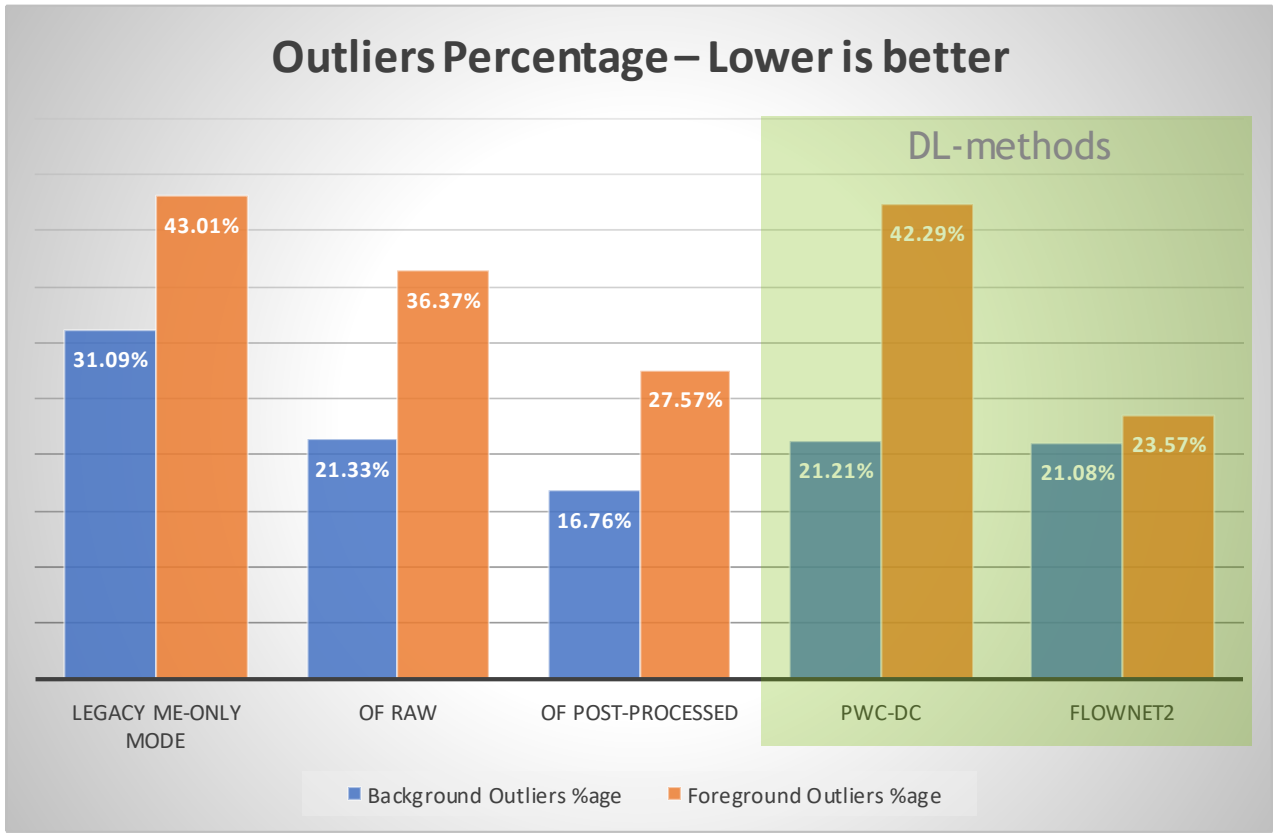

Outlier = Euclidian distance > 3 between OF vector and ground truth

OPTICAL FLOW QUALITY

Subjective Quality

- NVIDIA frame-rate-up-conversion
  - Video frame interpolation
  - ME-only mode (8×8), optical flow (4×4), optical flow with post-processing (1×1)
  - Subjective and objective quality comparison
- Results
  - Raw optical flow (4x4) based video interpolation better than ME-only mode (8x8) interpolation
  - Some video quality improvement with OF-post-processed (1x1) – content-dependent

VIDEO FRAME INTERPOLATION

Original 30 fps video

VIDEO FRAME INTERPOLATION

Upconverted 60 fps video

PERFORMANCE

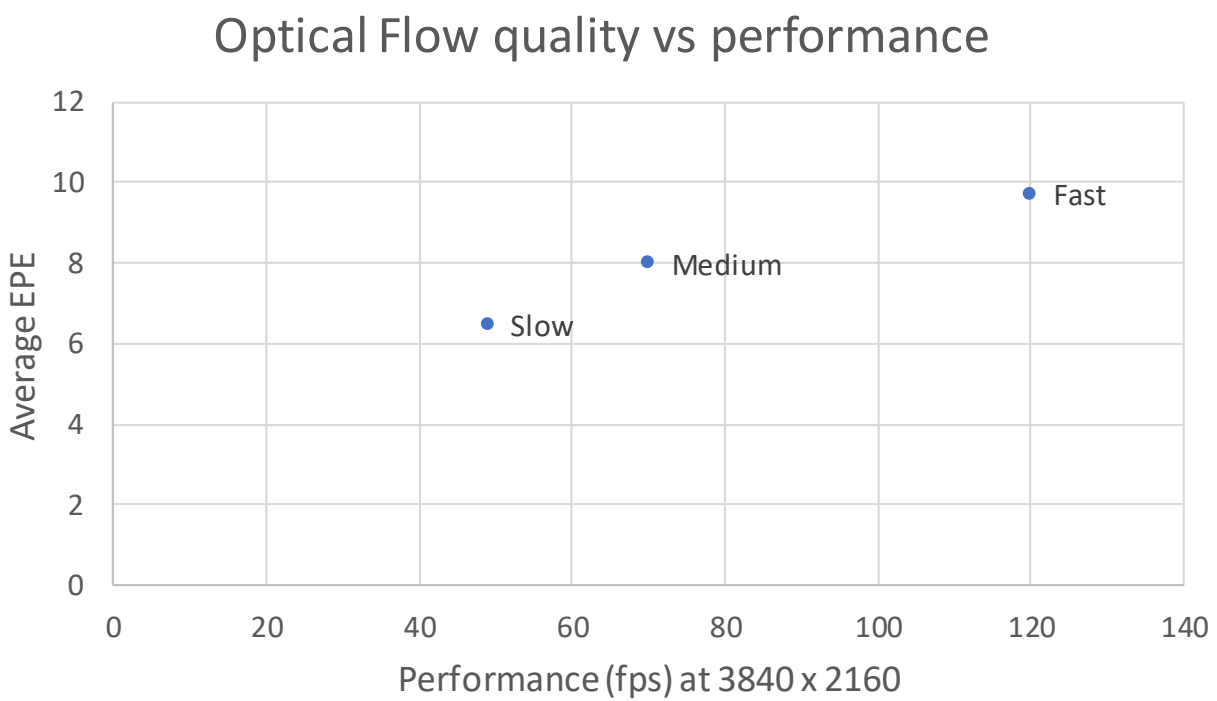

- 3 presets
- Fast/Medium – no CUDA processing
- Slow – pre/post-processing in CUDA
- Performance scales with resolution
- Cost calculation in CUDA (enable only if needed)


Nvidia's Optical Flow
NV OF time per frame: 2ms
GPU Utilization: 7%
Farneback's Optical Flow
Farneback's OF time per frame: 8ms
GPU Utilization: 13%
Lucas Kannade's Optical Flow
LK's OF time per frame: 25ms
GPU Utilization: 10%
Vchill
This PC
Recycle Bin
GeForce Experience
Statistics

END-TO-END APPLICATIONS

END-TO-END USE-CASES

Applications

- Video comprehension/classification
  - 2x better accuracy compared to no optical flow with UCF-101
  - Makes OF-assisted-video-comprehension usable
- Optical-flow-assisted video inter/extrapolation
  - Objective and subjective quality comparable to FlowNet2
  - Turing enables *real-time* optical-flow-assisted video interpolation

OPTICAL FLOW-ASSISTED VIDEO CLASSIFICATION

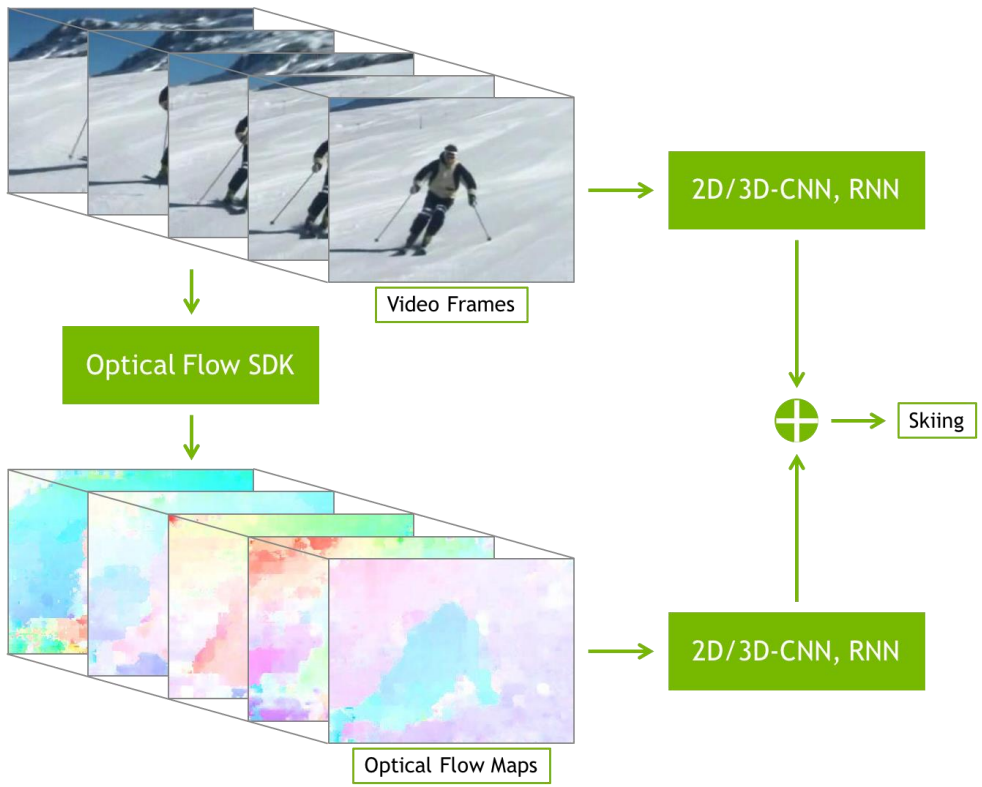

VIDEO CLASSIFICATION

Enables world class classification accuracy with real time performance

Turing hardware :
→ Optical Flow reduces error rates by 2x
→ 20+ streams 720p inference

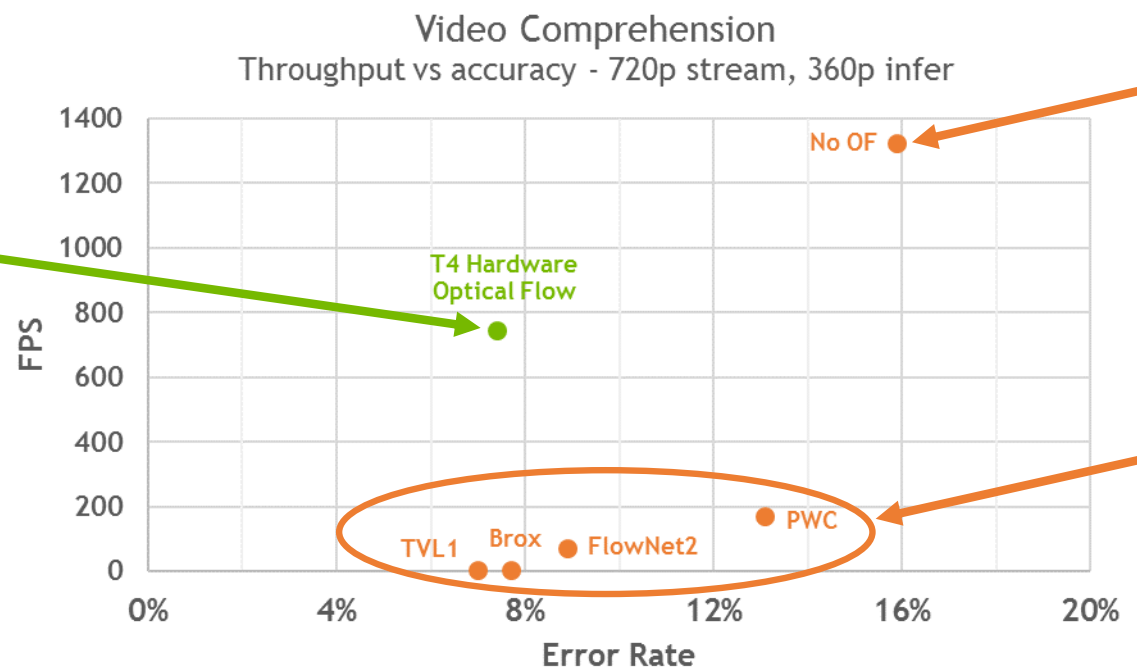

Image only video classification has high error rates

Optical flow significantly reduces error rates, but DL based OF is unusably slow

TURING OPTICAL FLOW

High quality video frame interpolation at 4K in real-time

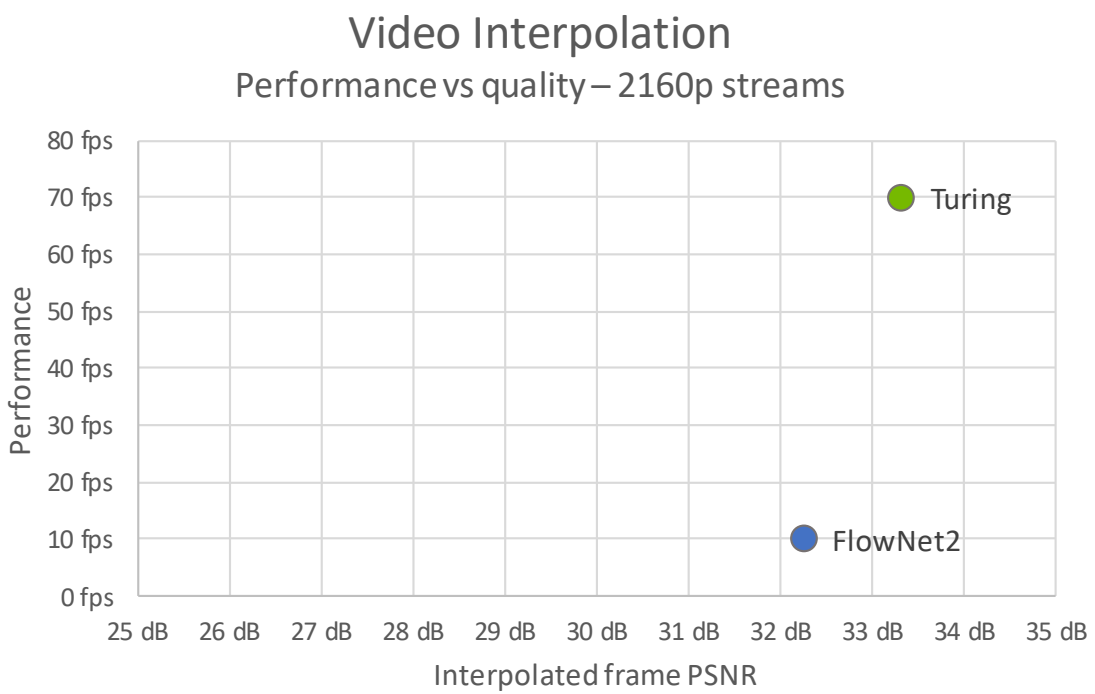

Video Interpolation

- 60 fps ➔ 120 fps at 4K in real-time
- 7x perf vs FlowNet2
- 1 dB better objective quality (PSNR) than FlowNet2-assisted interpolation
- Similar visual quality as FlowNet2-assisted interpolation

ROADMAP

ROADMAP

Optical Flow SDK 1.1

- Q3 2018
- Improved quality via post-processing
- 1x1 flow vectors
- Integration into DALI, Pytorch and other DL frameworks

RESOURCES

Optical Flow SDK: https://developer.nvidia.com/opticalflow-sdk

Support: video-devtech-support@nvidia.com

Video & Optical Flow SDK forums:
https://devtalk.nvidia.com/default/board/175/video-technologies/

Connect with Experts (CE9103): Wednesday, March 20, 2019, 3:00 pm

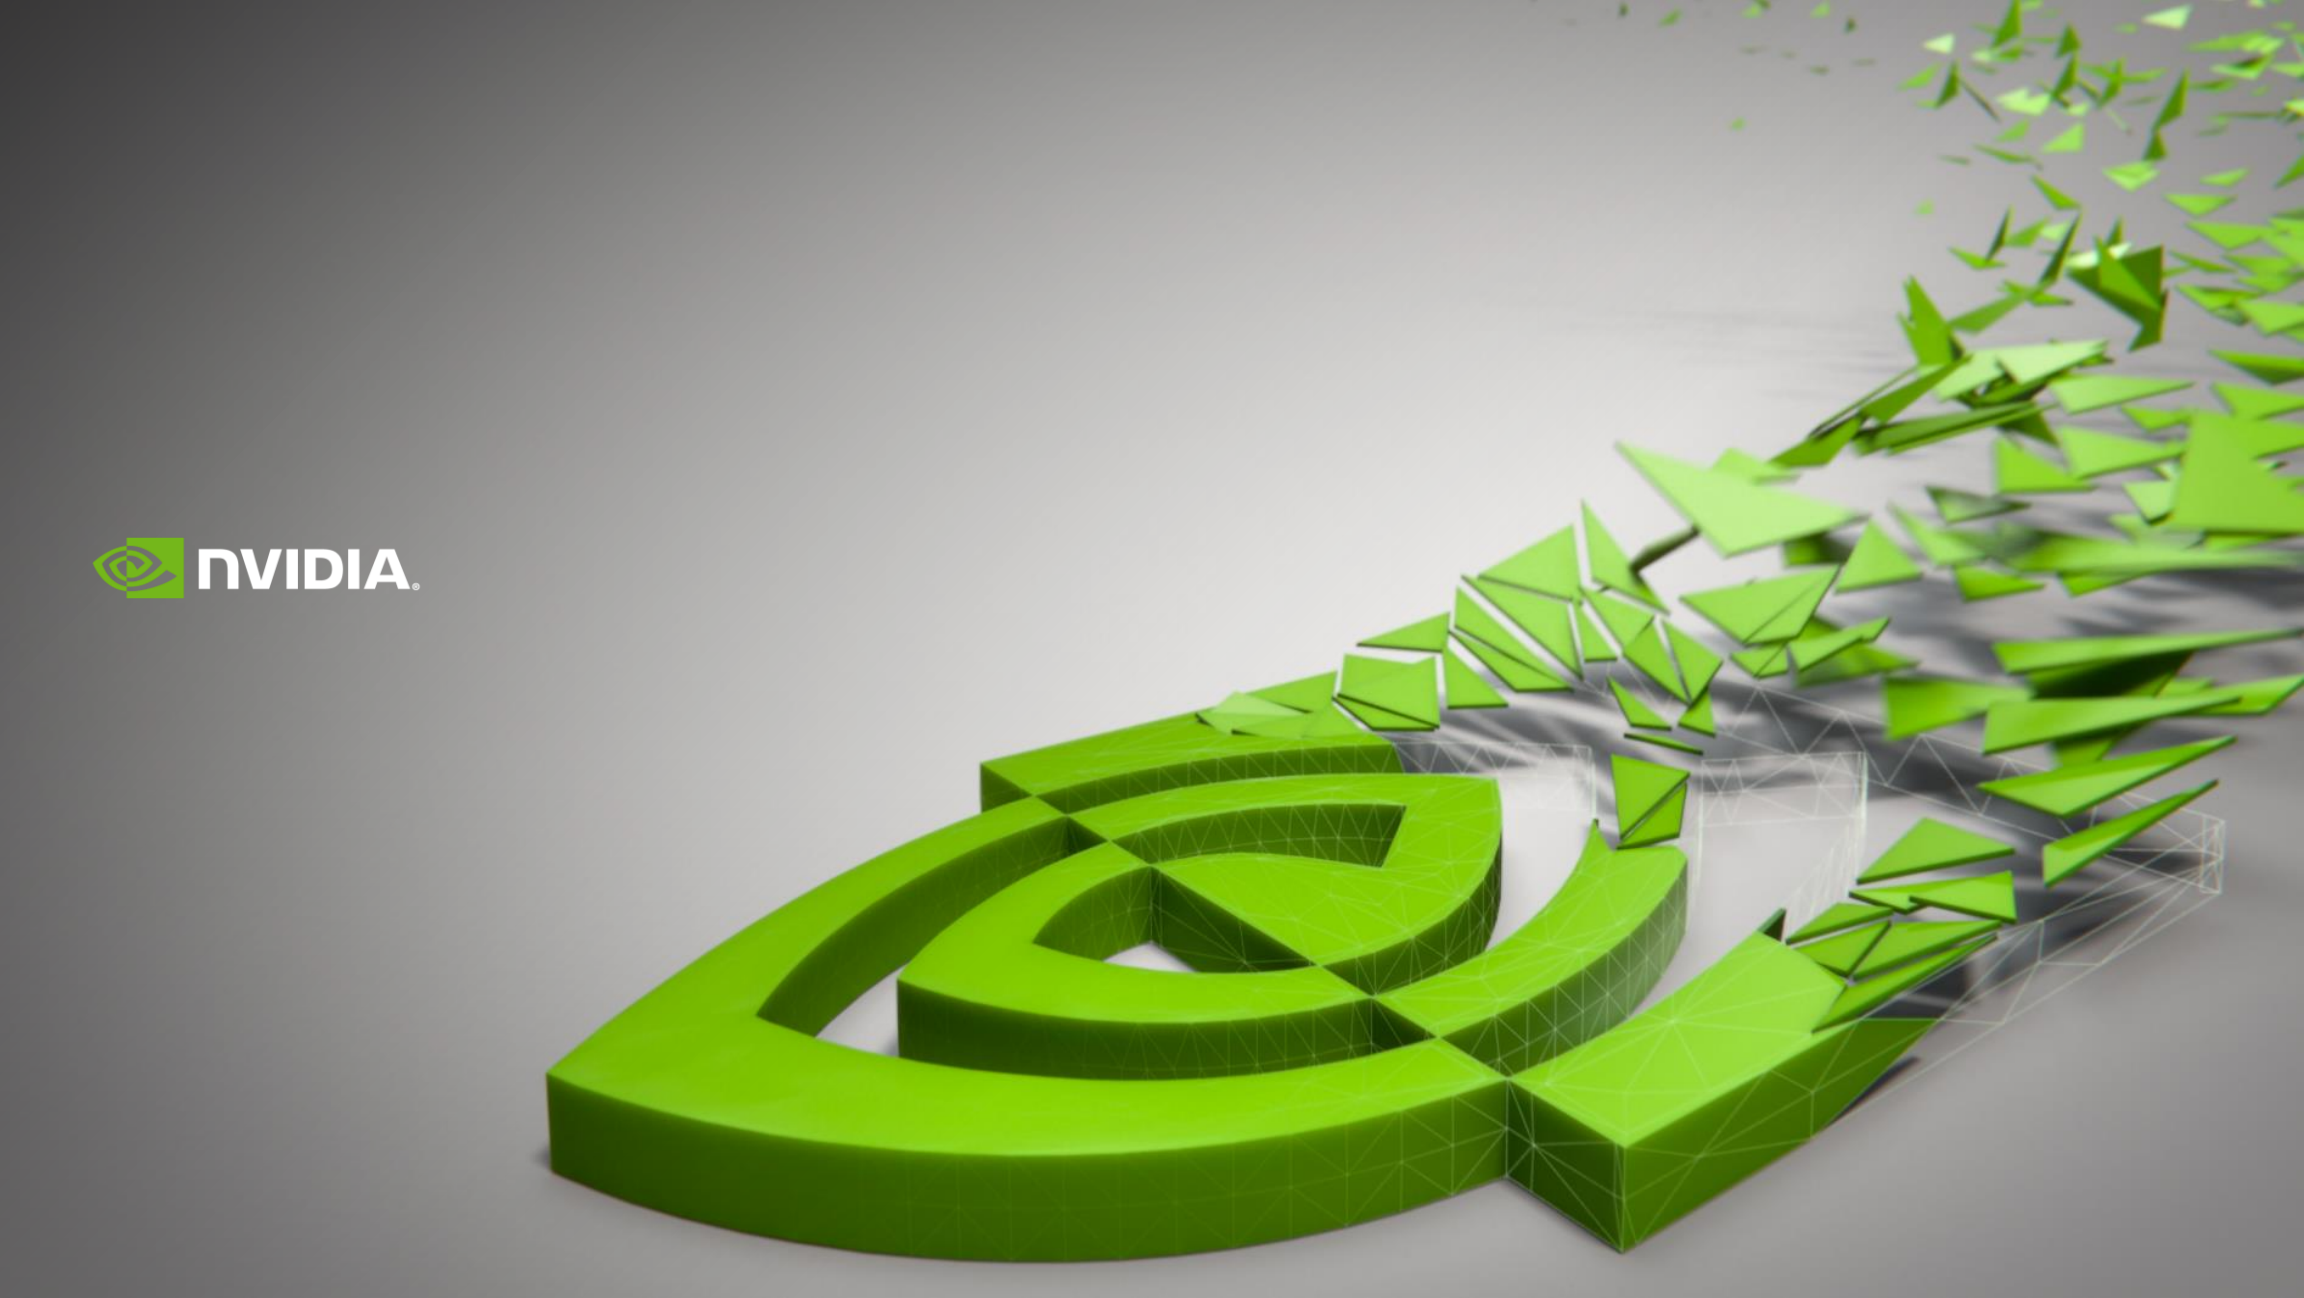
NVIDIA