GTC 2013 | March 18-21 | San Jose, CA

The Smartest People. The Best Ideas. The Biggest Opportunities.

Opportunities for Participation:

SPEAK - Showcase your work among the elite of graphics computing

- Call for Sessions: August 2012
- Call for Posters: October 2012

REGISTER - learn from the experts and network with your peers

- Use promo code **GM10SIGG** for a 10% discount

SPONSOR - Reach influential IT decision-makers

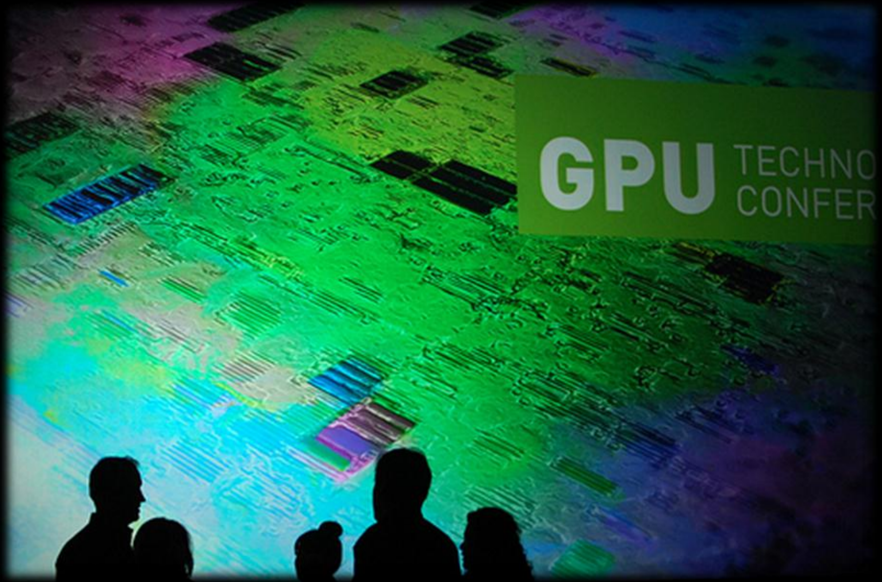

Learn more at www.gputechconf.com

Mobile Visual Computing

Kari Pulli
Senior Director, Research

NVIDIA

Overview

- Tegra – NVIDIA's mobile SOC (system-on-chip)
- OpenCV – a de-facto standard API for computer vision
- FCam – camera control for computational photography
- CUDA – flexible image processing with Carma

ARM Enables Energy Efficient Computing

Personal Computing

ARM Servers

ARM is Pervasive and Open

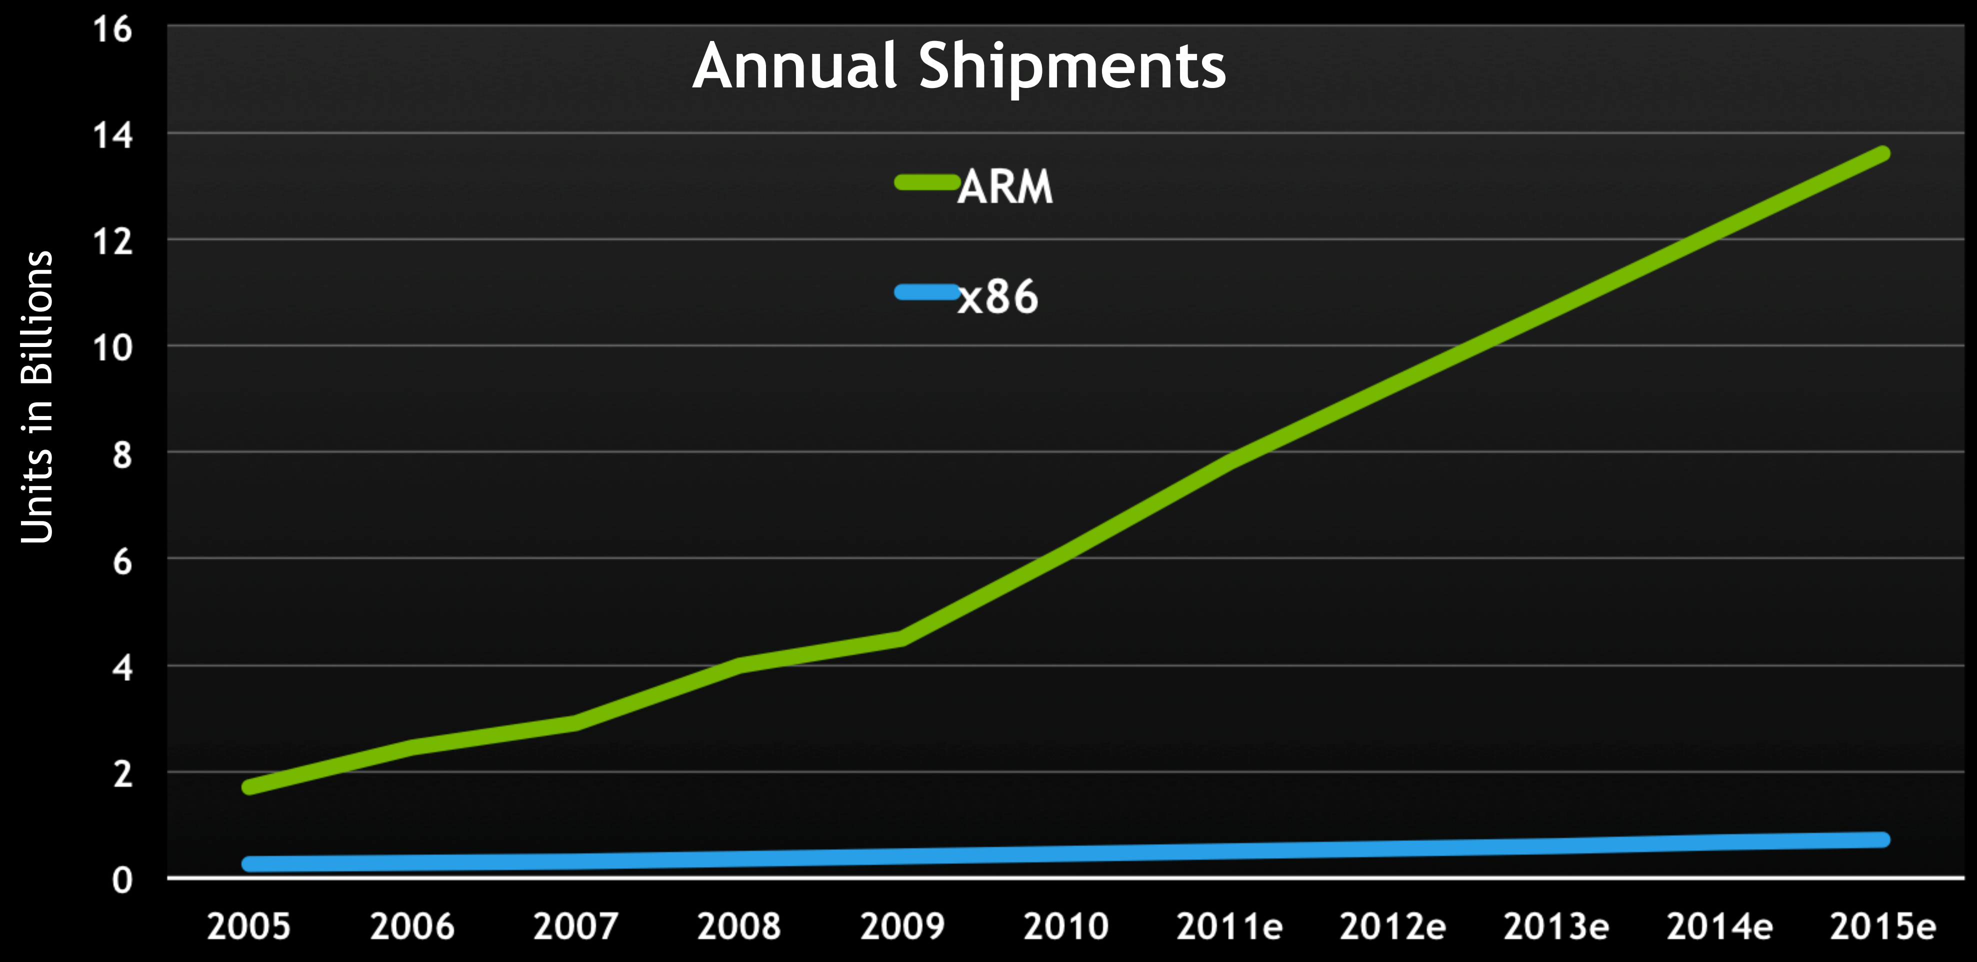

Source: ARM, Mercury Research, NVIDIA

Tegra 3 SOC

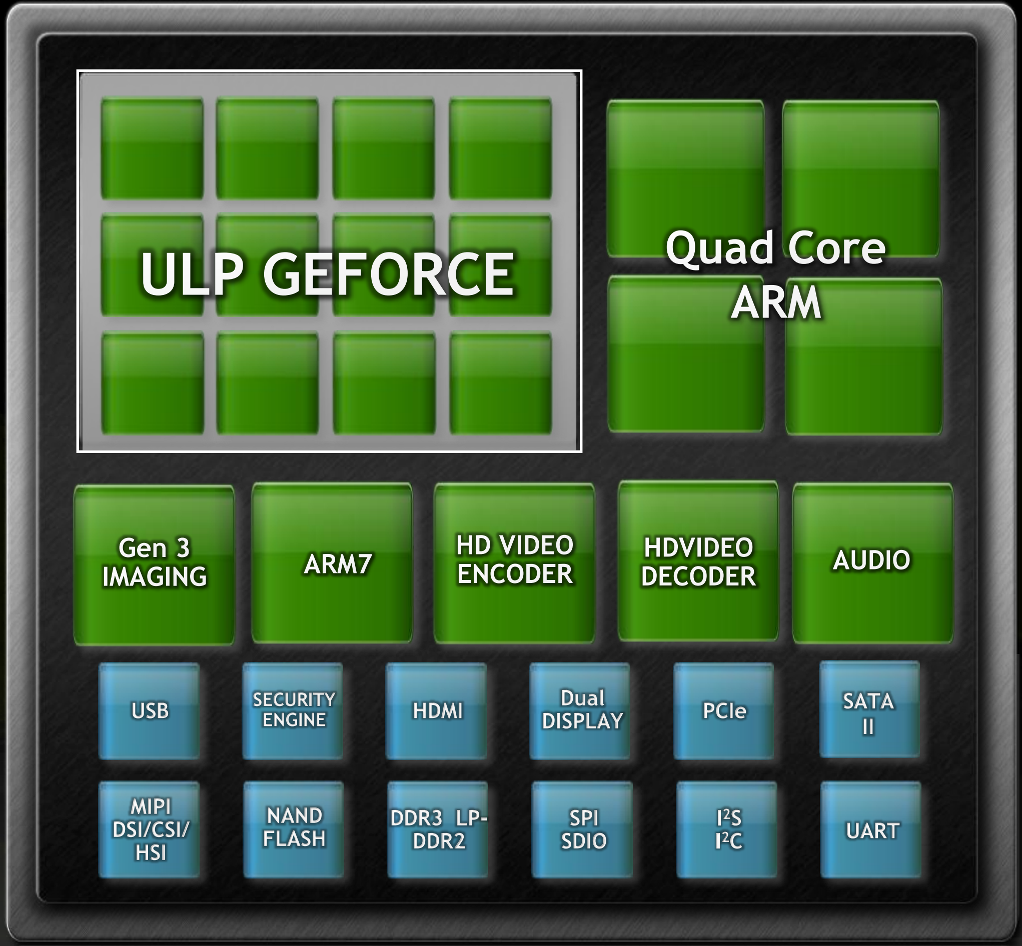

CPU	3X Performance *Quad Core NEON*
POWER	20x Lower Power *Due to ULP mode*
VIDEO	4X Complexity *1080i/p High Profile*
GRAPHICS	3X Performance *12 Core, Dual Pixel Pipe*
MEMORY	3X bandwidth *DDR3L up to 1600 data rate*
IMAGING	Better noise reduction & color rendition *Two simultaneous streams*
AUDIO	HD Audio *7.1 channel surround*
STORAGE	2 - 6X faster *e.MMC 4.4 and SATA-II*

4 + 1

- An extra core
  - “companion core”
  - “shadow core”
  - built with a low-leakage process
    - lower peak performance
- All cores identical
  - the same code runs on all cores
  - application programmer doesn’t have to worry
    - power is adapted automagically

5 CPU Cores

Most Common Use Cases

Gaming, Web, Email

~80%

~20%

Active Standby

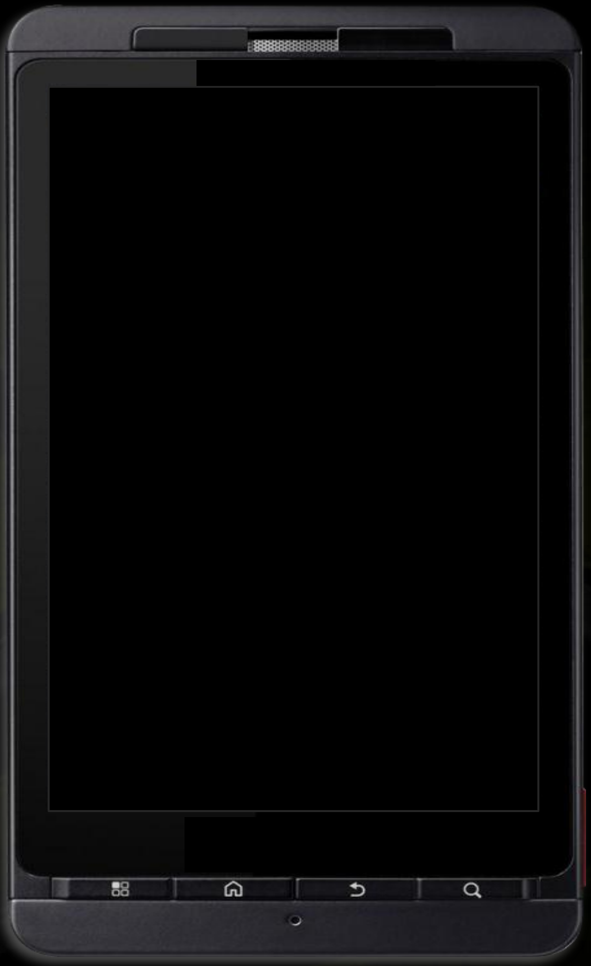

What's really running

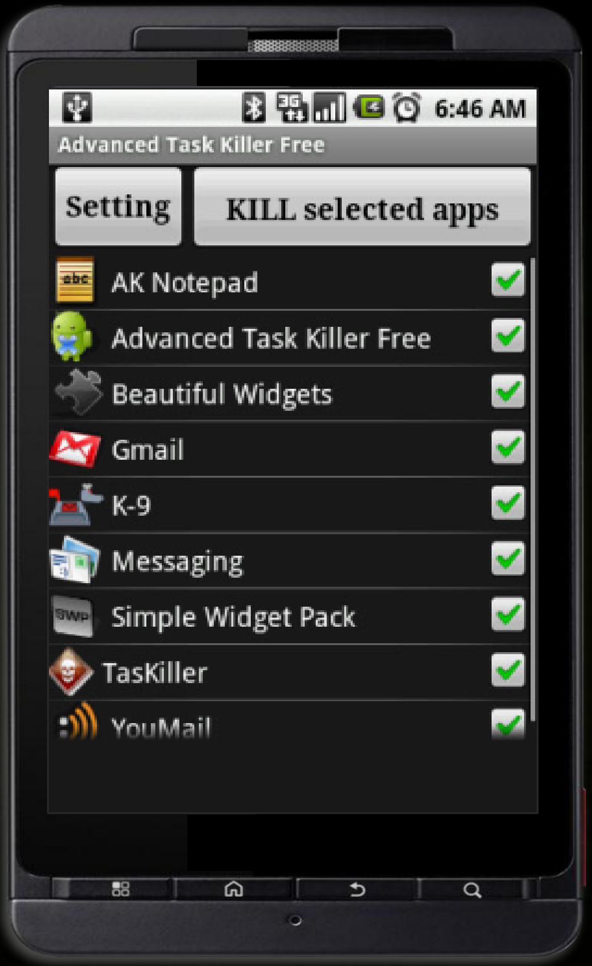

Variable SMP

- ✓ “Companion” Core Active
- ✓ 0 to 500 MHz
- ✓ 0 to 200 mW

ARM CPU Cores

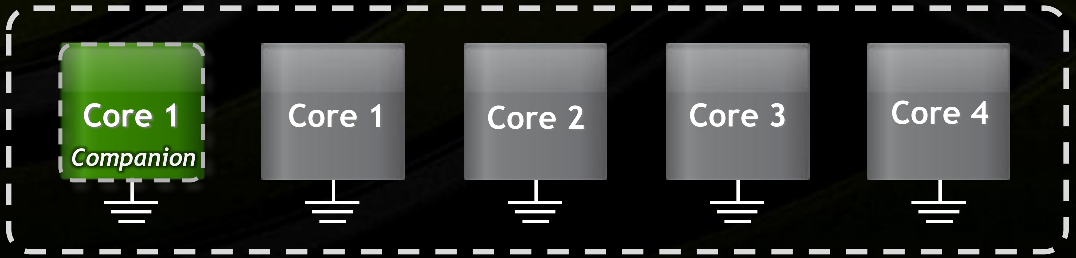

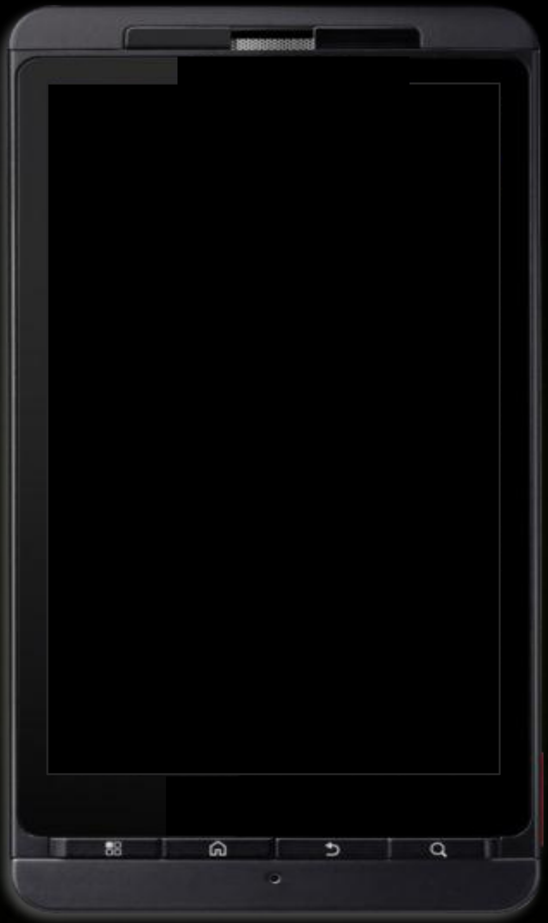

Active Standby,
Music, Video

Variable SMP

- ✓ 1 CPU Core Active
- ✓ 0 to 1.7 GHz
- ✓ 30 to 400 mW

ARM CPU Cores

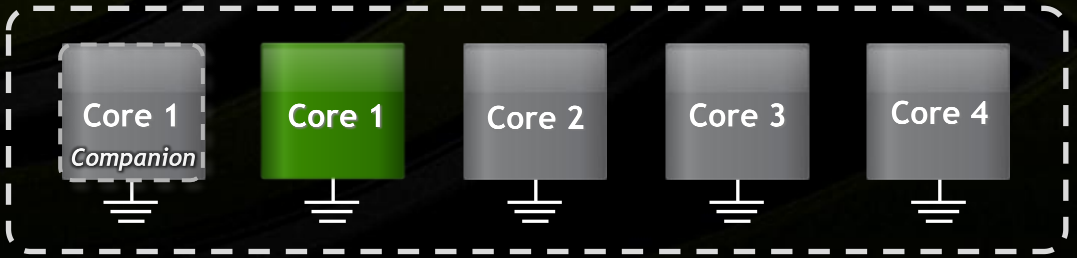

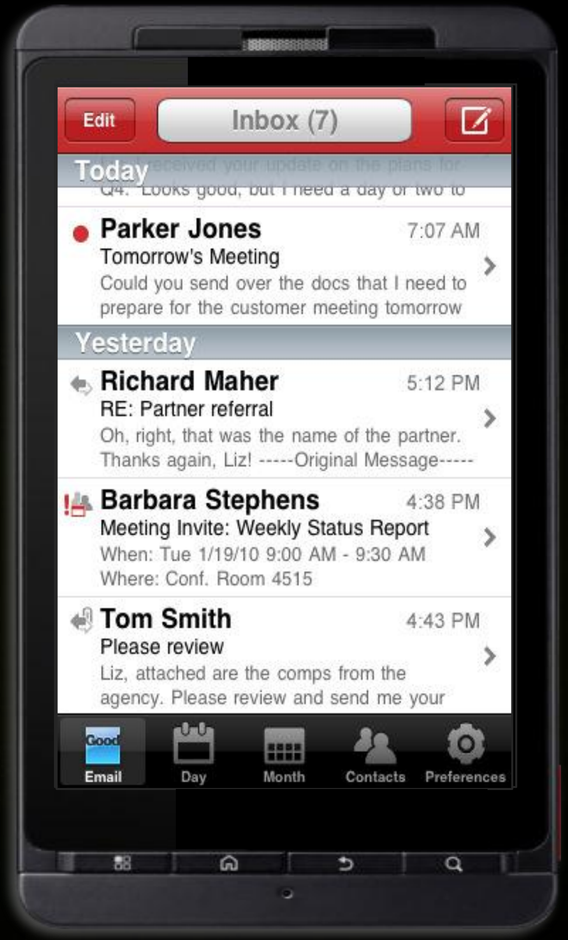

Email

Variable SMP

- ✓ 2 CPU Cores Active
- ✓ 0 to 1.6 GHz
- ✓ 50 to 800 mW

ARM CPU Cores

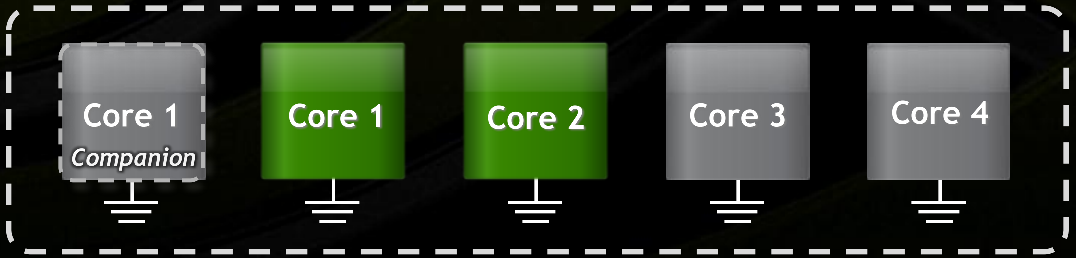

Web

Variable SMP

- ✓ 4 CPU Cores Active
- ✓ 0 to 1.6 GHz
- ✓ 100 to 1600 mW

ARM CPU Cores

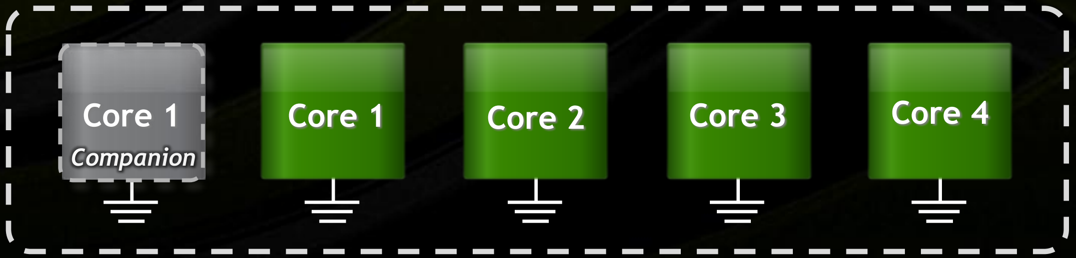

Gaming, Multimedia Apps

OpenCV

- For the computation part of computer vision and computational photography
  - a de-facto standard
  - optimized for Tegra (4x[ARM+NEON] + GPU)
  - OpenCV4Tegra (version 2.4.2) in Tegra Android Dev. Prog. in August

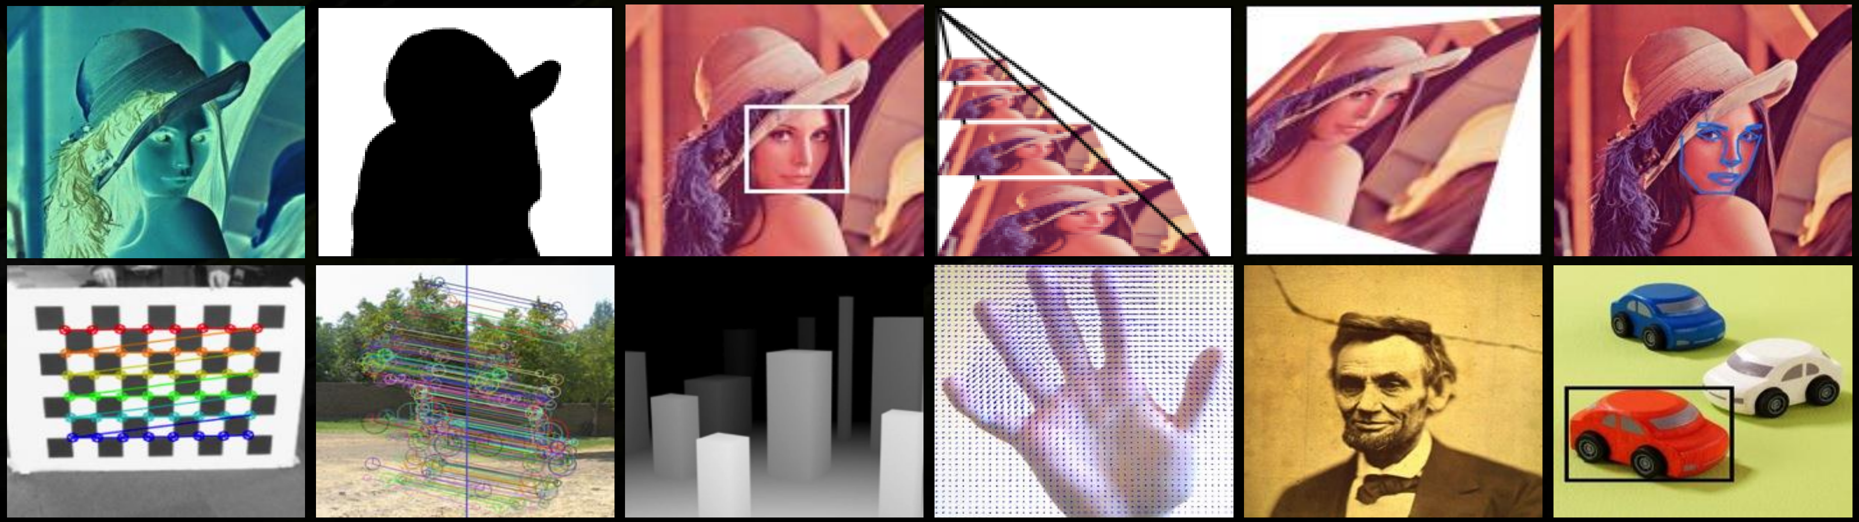

OpenCV examples

- Stereo
- Tracking for Augmented Reality

Stereo matching on Tegra Dev Board (Cardhu)

Calibration

Rectification

Matching

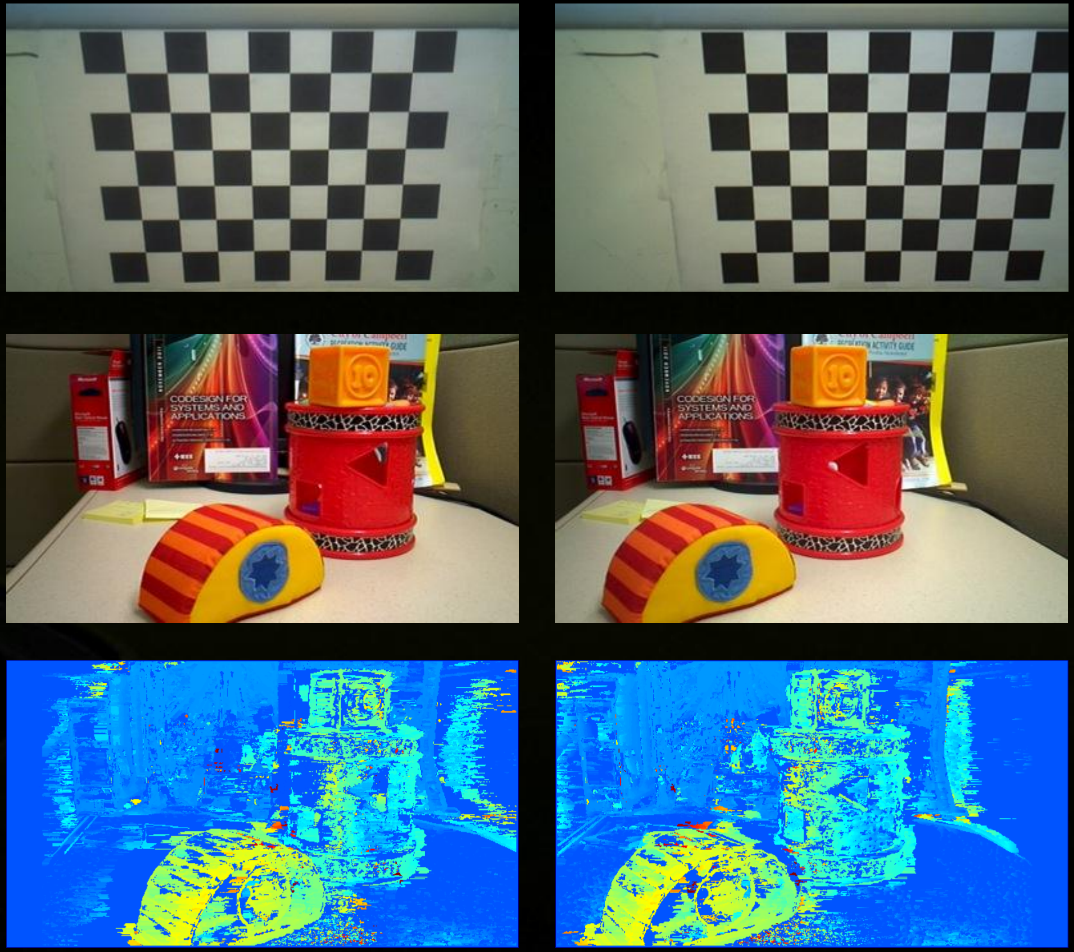

Flash stereo

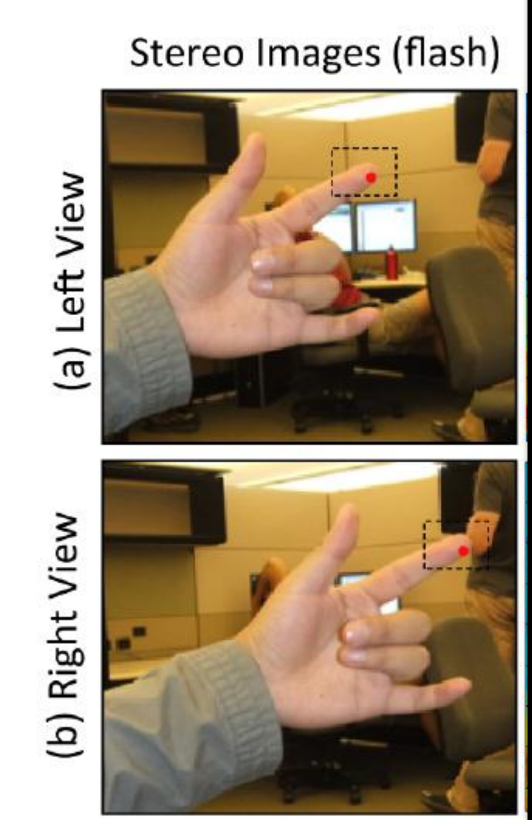

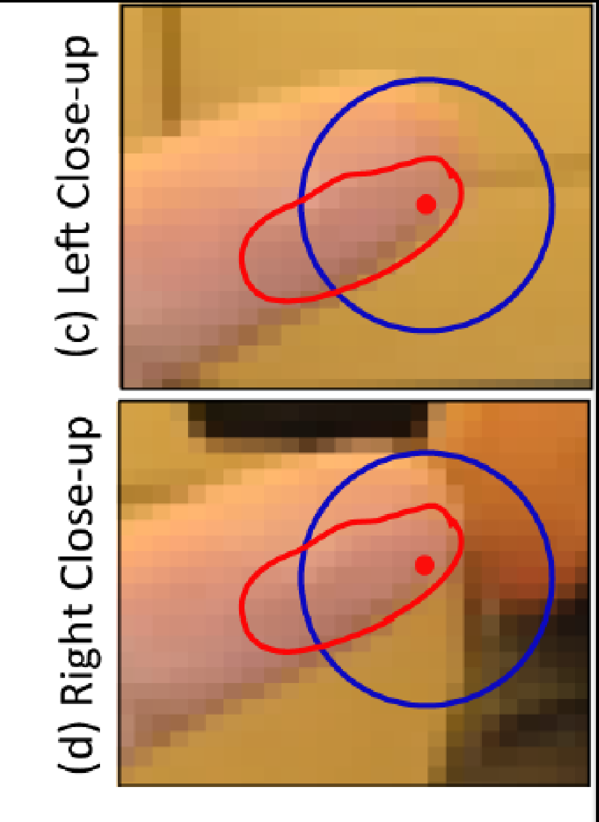

Augmented Reality (AR)

- What
  - visual tracking
  - of some known environment
  - overlay graphics

- Conventional AR
  - track markers
  - track a picture

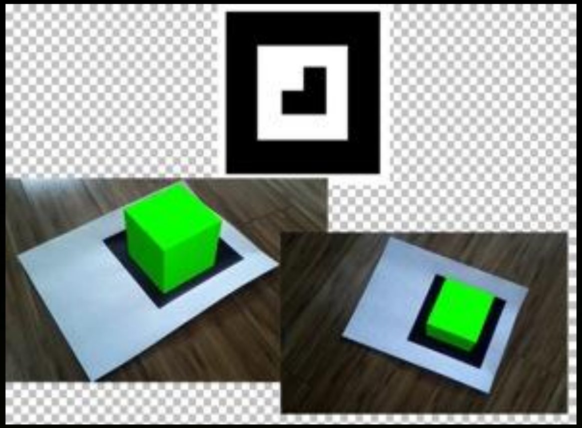

AR game on Tegra (2009)

Augmented Reality (AR) with SLAM

- **SLAM: build the model as you go**
  - **Simultaneous Localization and Mapping**
  - **Examples of Visual SLAM**
    - **PTAM (2007, 2009)**
    - **DTAM (2011)**
- **Overview of visual SLAM for AR**
  - **Simultaneous tracking and mapping**

PTAM (2007)

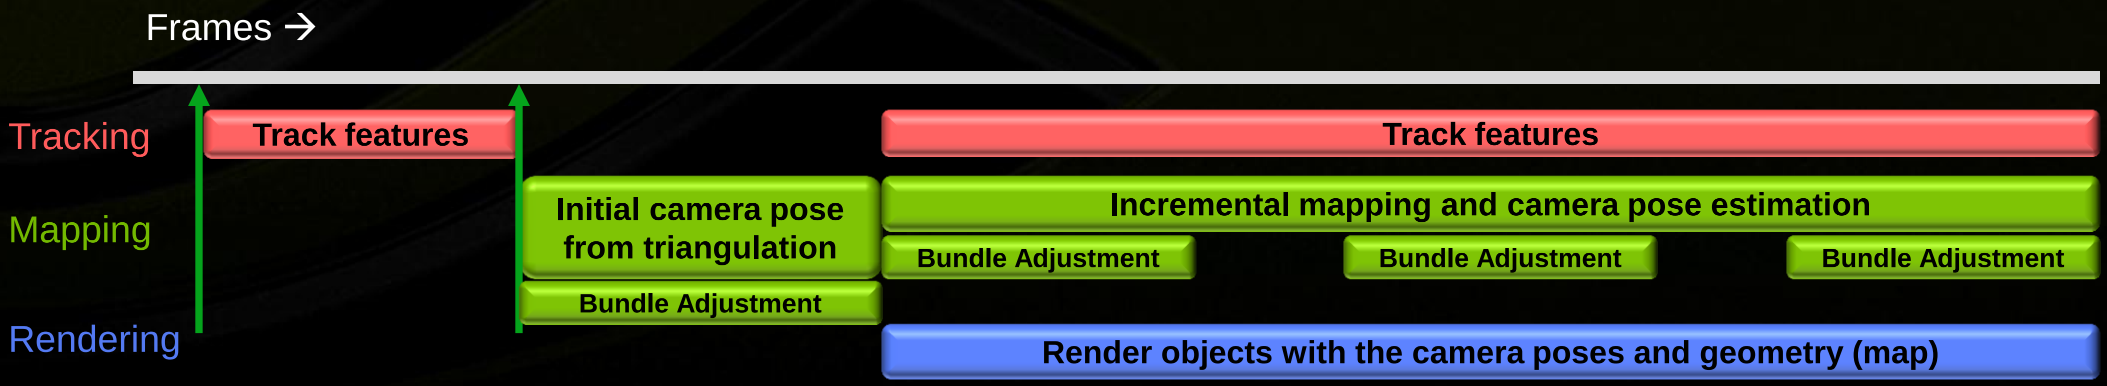

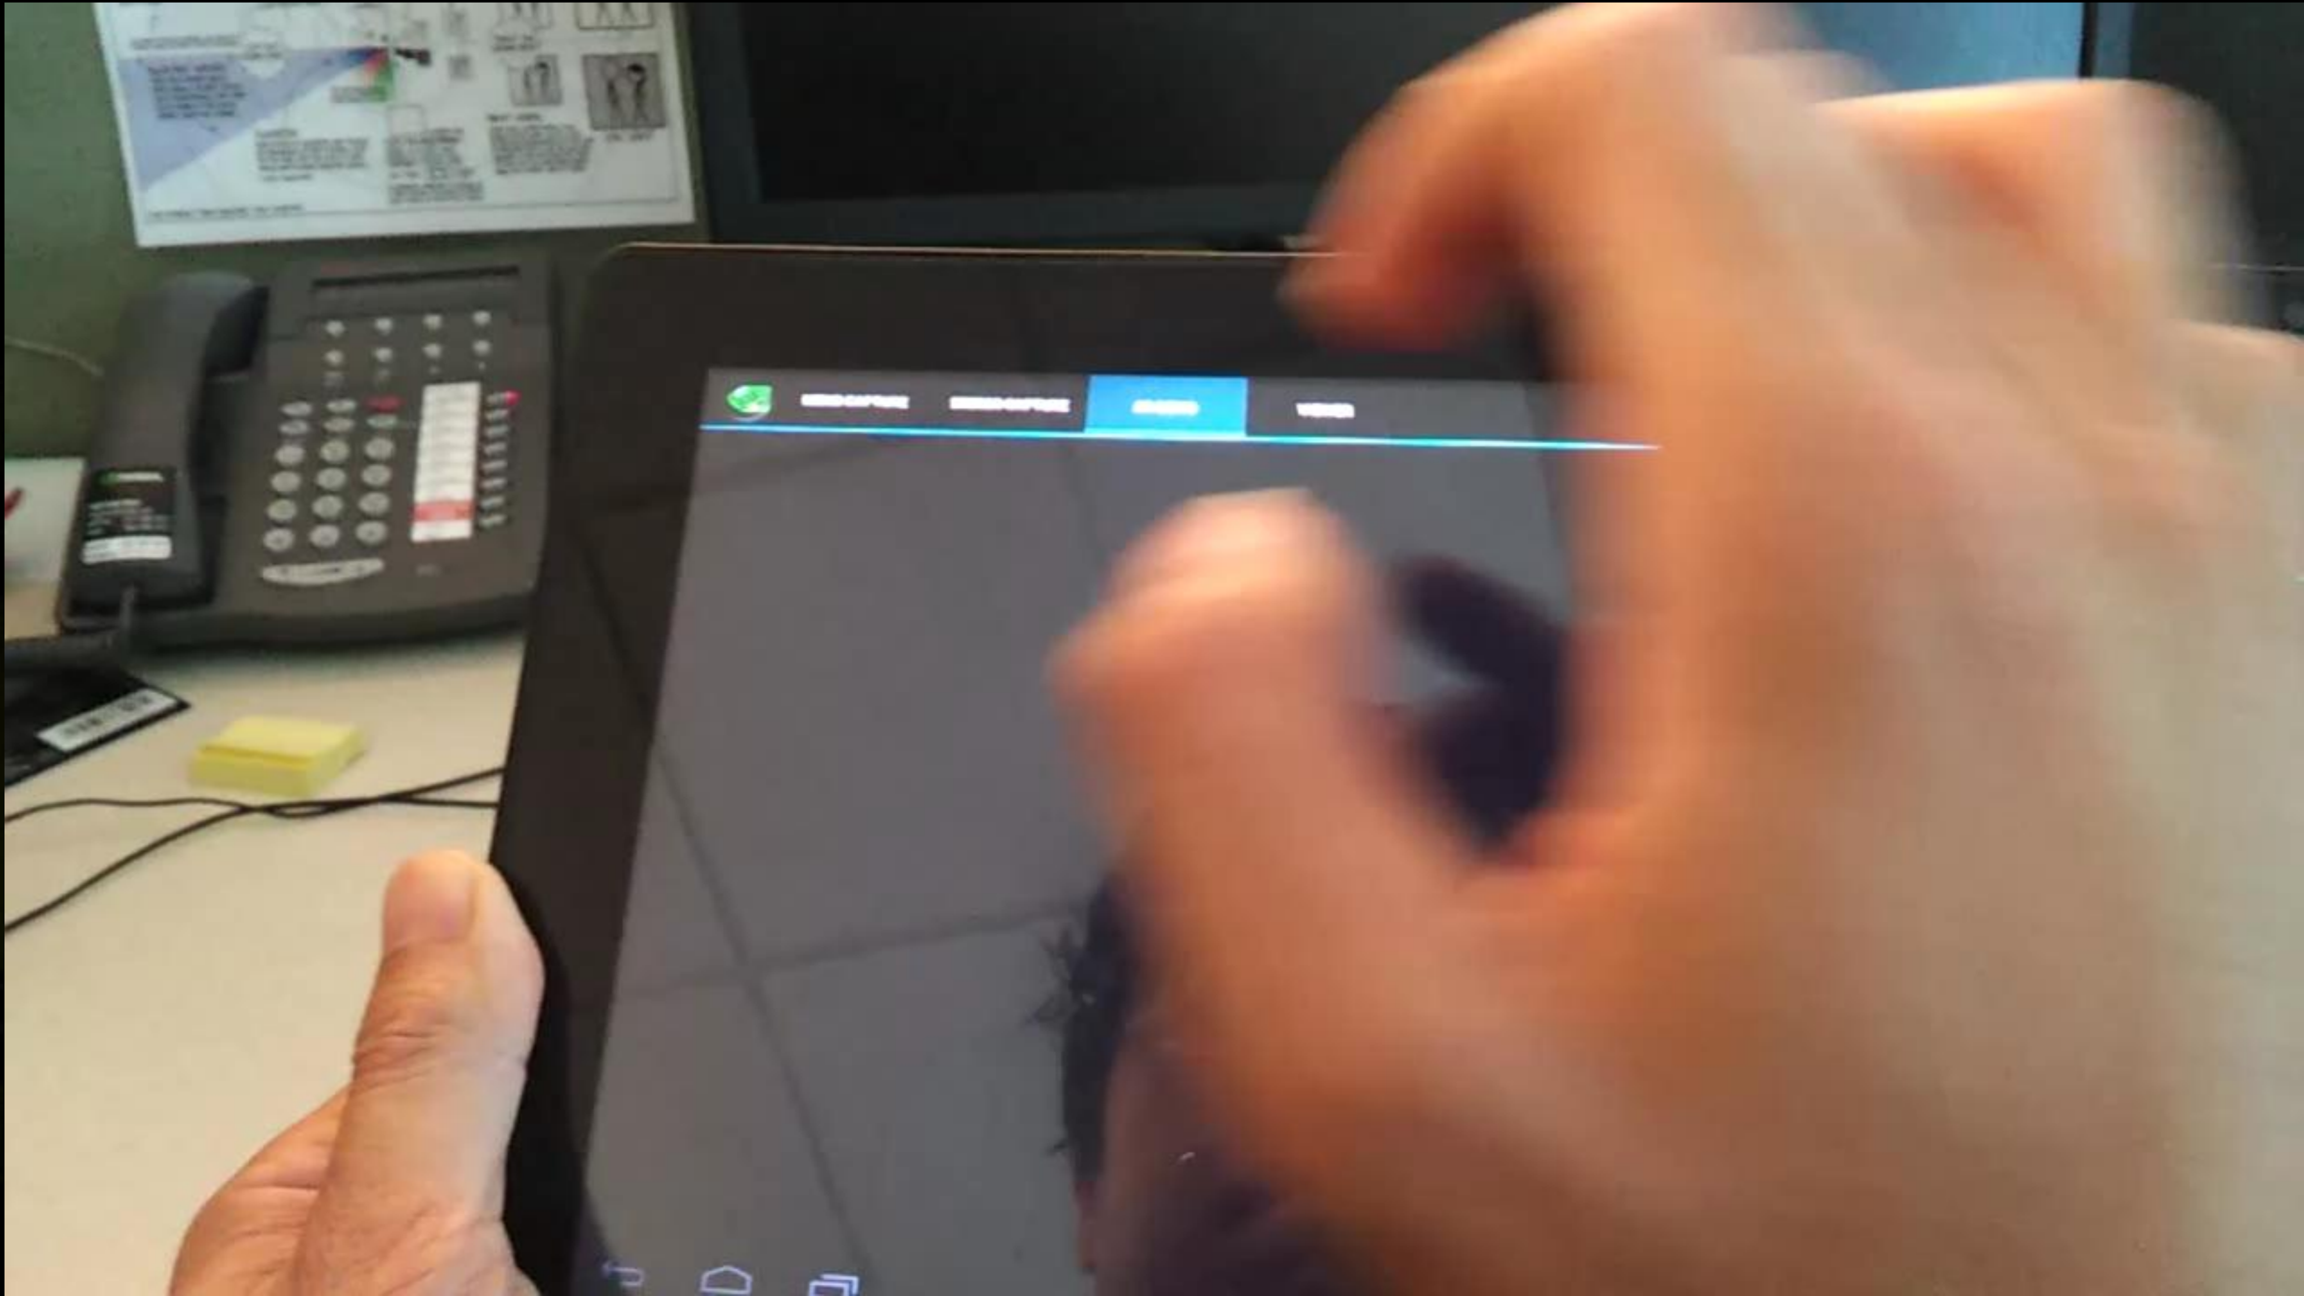

FCam

- Open Source API for camera control for Computational Photography
  - Nokia Linux phones
  - NVIDIA Tegra 3 dev boards

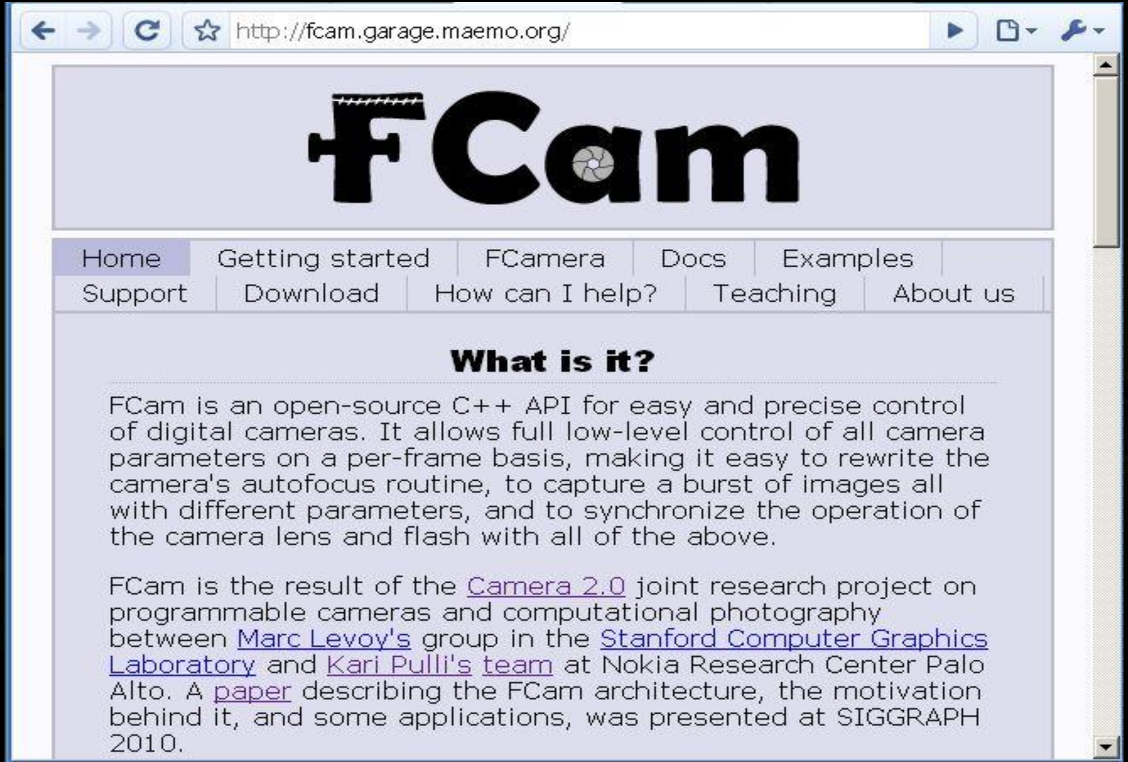

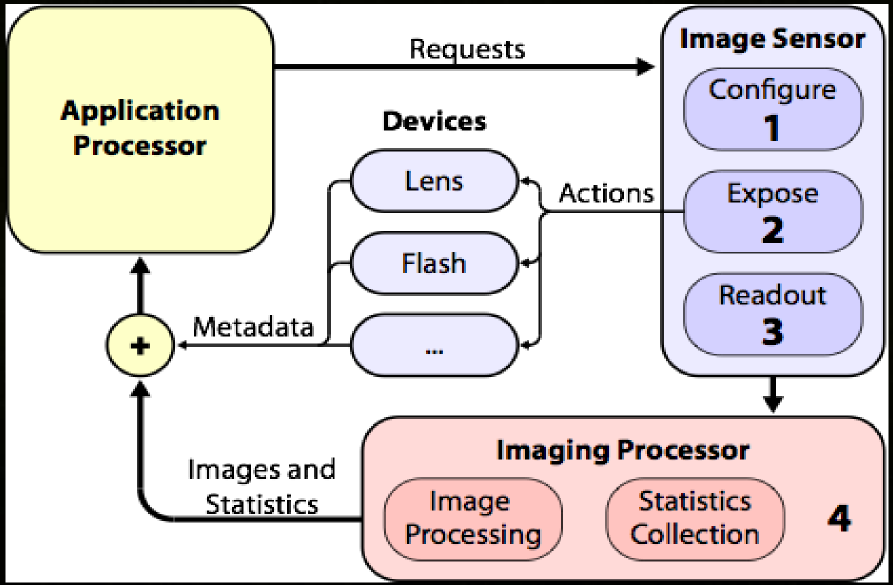

Traditional sensor model does not work for Comp. Photography

- Real image sensors are pipelined
  - while one frame exposing
  - next one is being prepared
  - previous one is being read out
- Viewfinding / video mode:
  - pipelined, high frame rate
  - settings changes take effect sometime later
- Still capture mode:
  - need to know which parameters were used
  - → reset pipeline between shots → slow

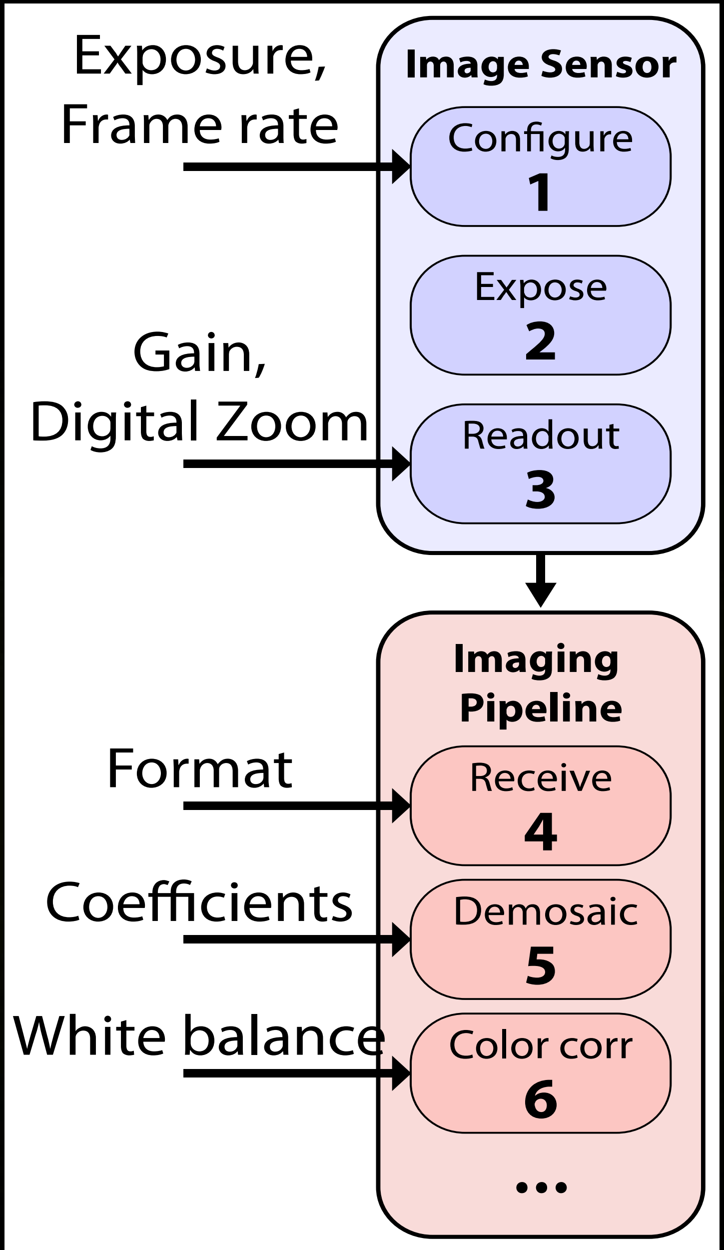

The FCam Architecture

- No global state
  - instead, state travels with image requests
  - every stage in pipeline may have different state
  - → allows deterministic, fast state changes
- Synchronize devices
  - flash
  - lens
  - capture sound
  - gyro
  - …

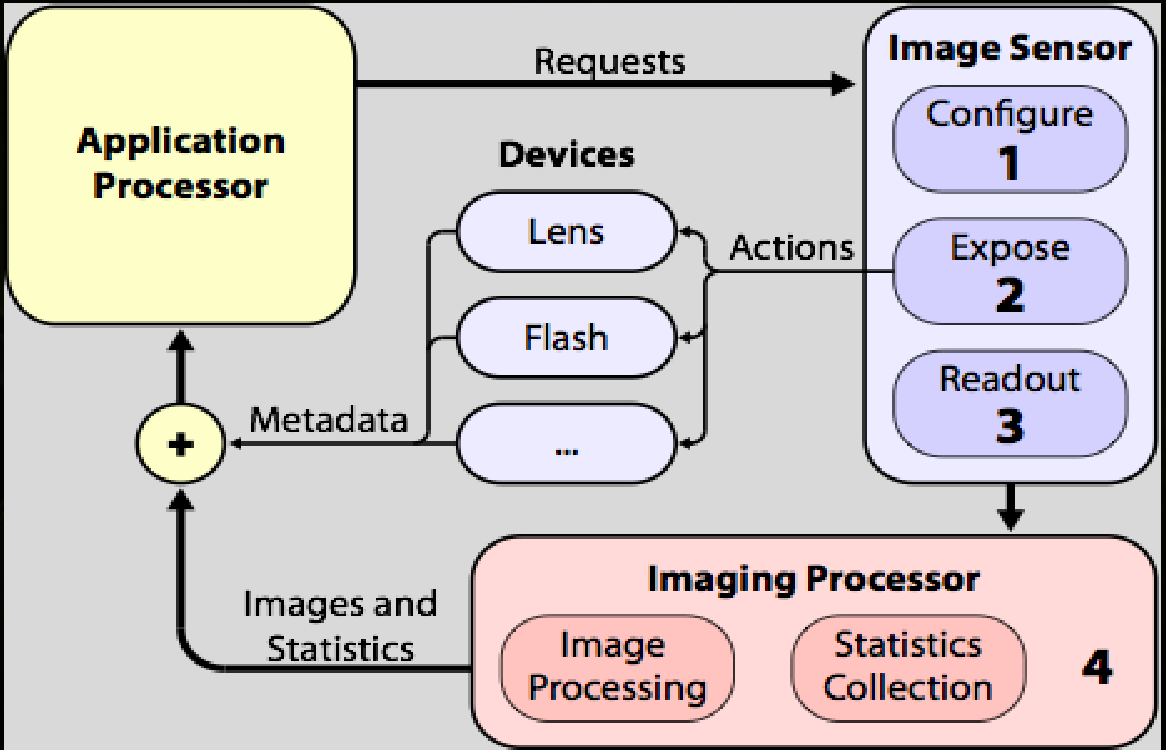

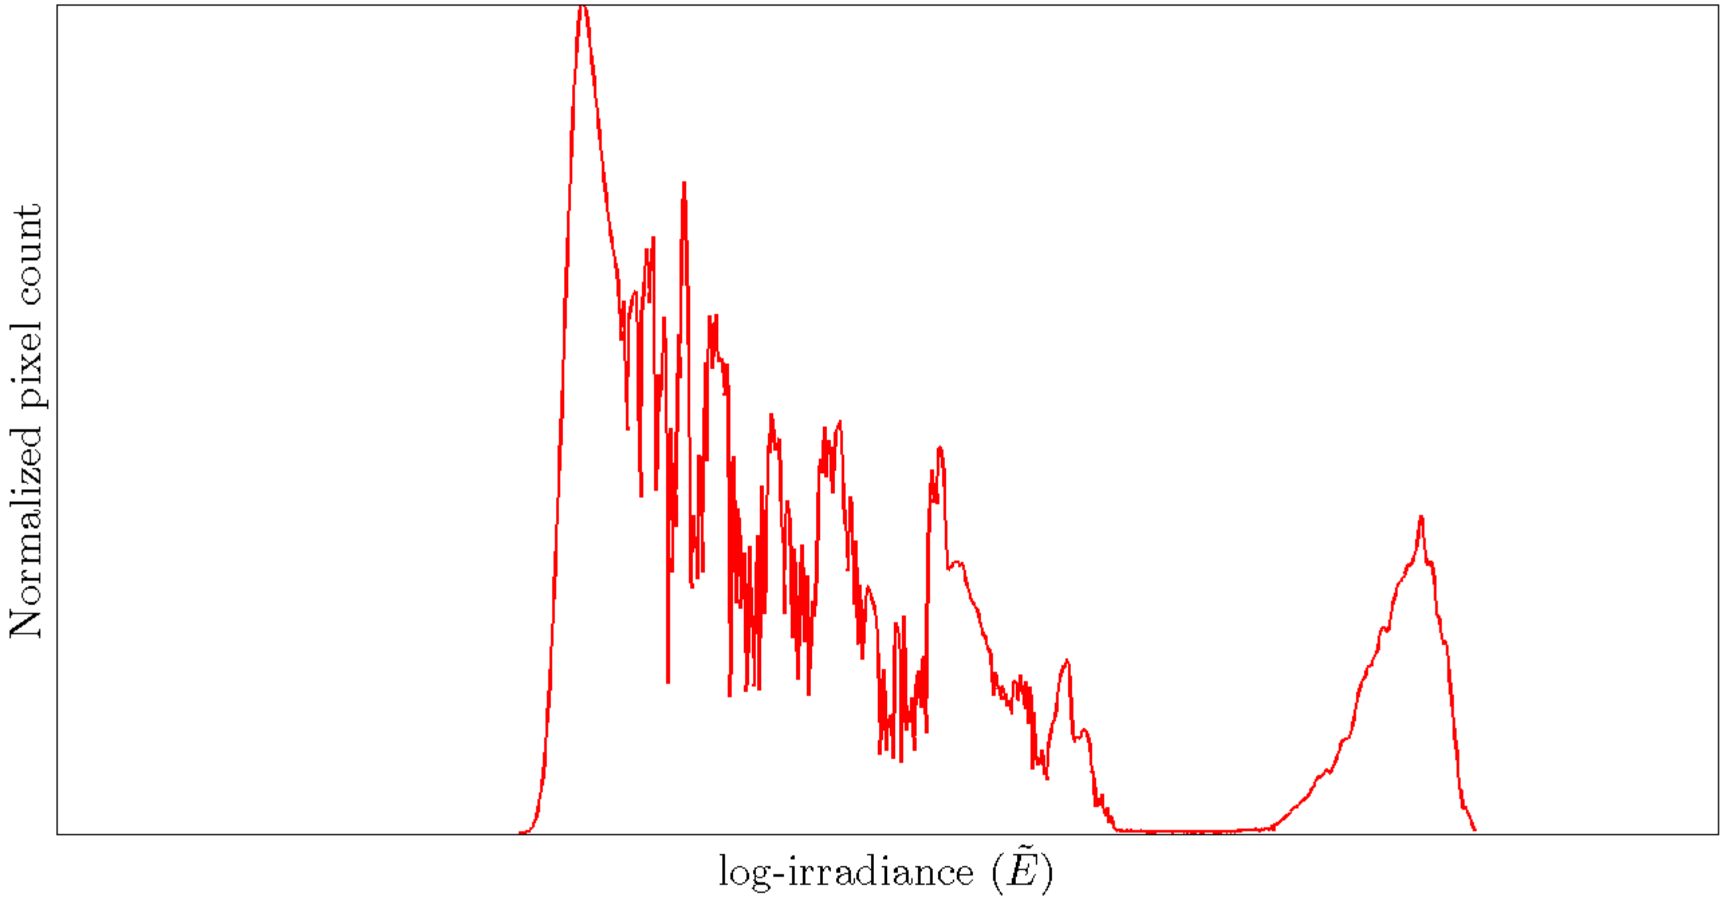

Normalized pixel count
log-irradiance ($\tilde{E}$)

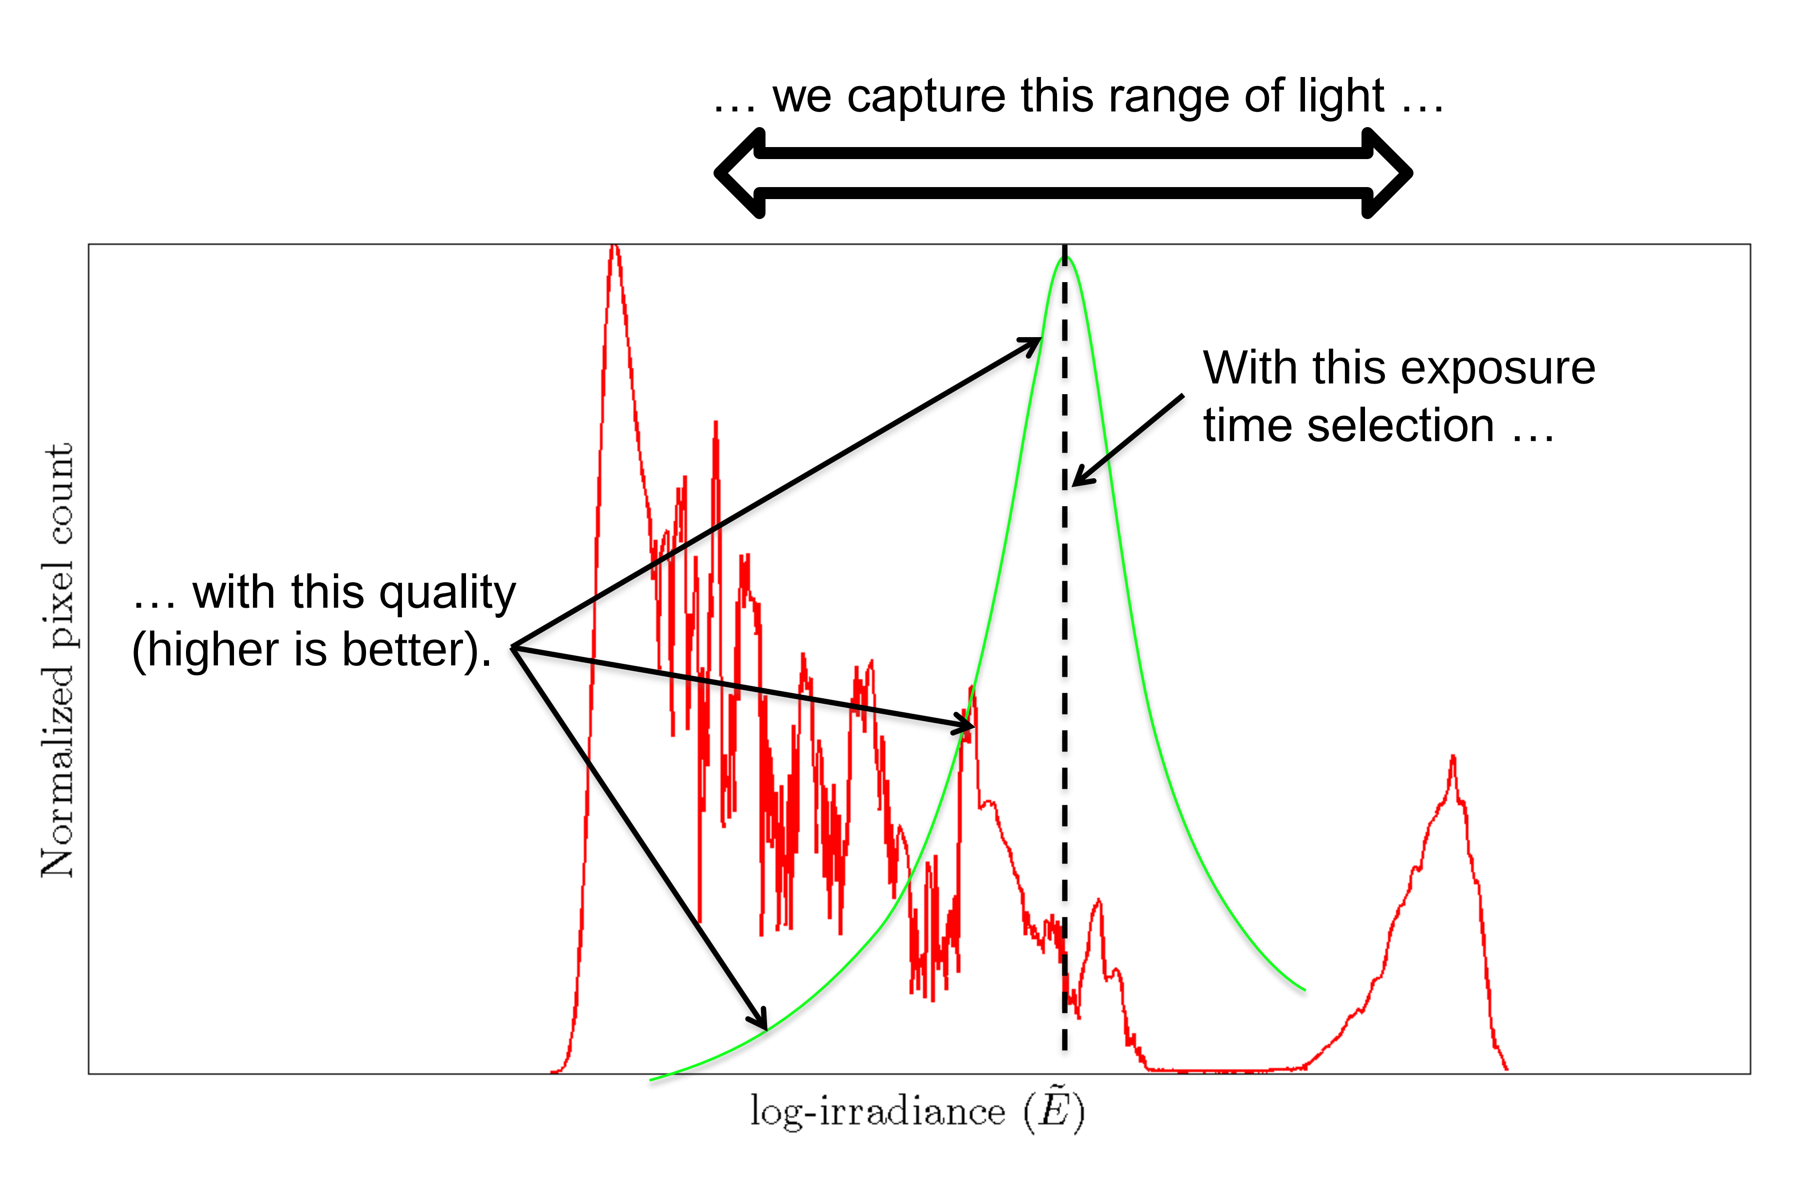
… we capture this range of light …
With this exposure
time selection …
… with this quality
(higher is better).
Normalized pixel count
log-irradiance ($\tilde{E}$)

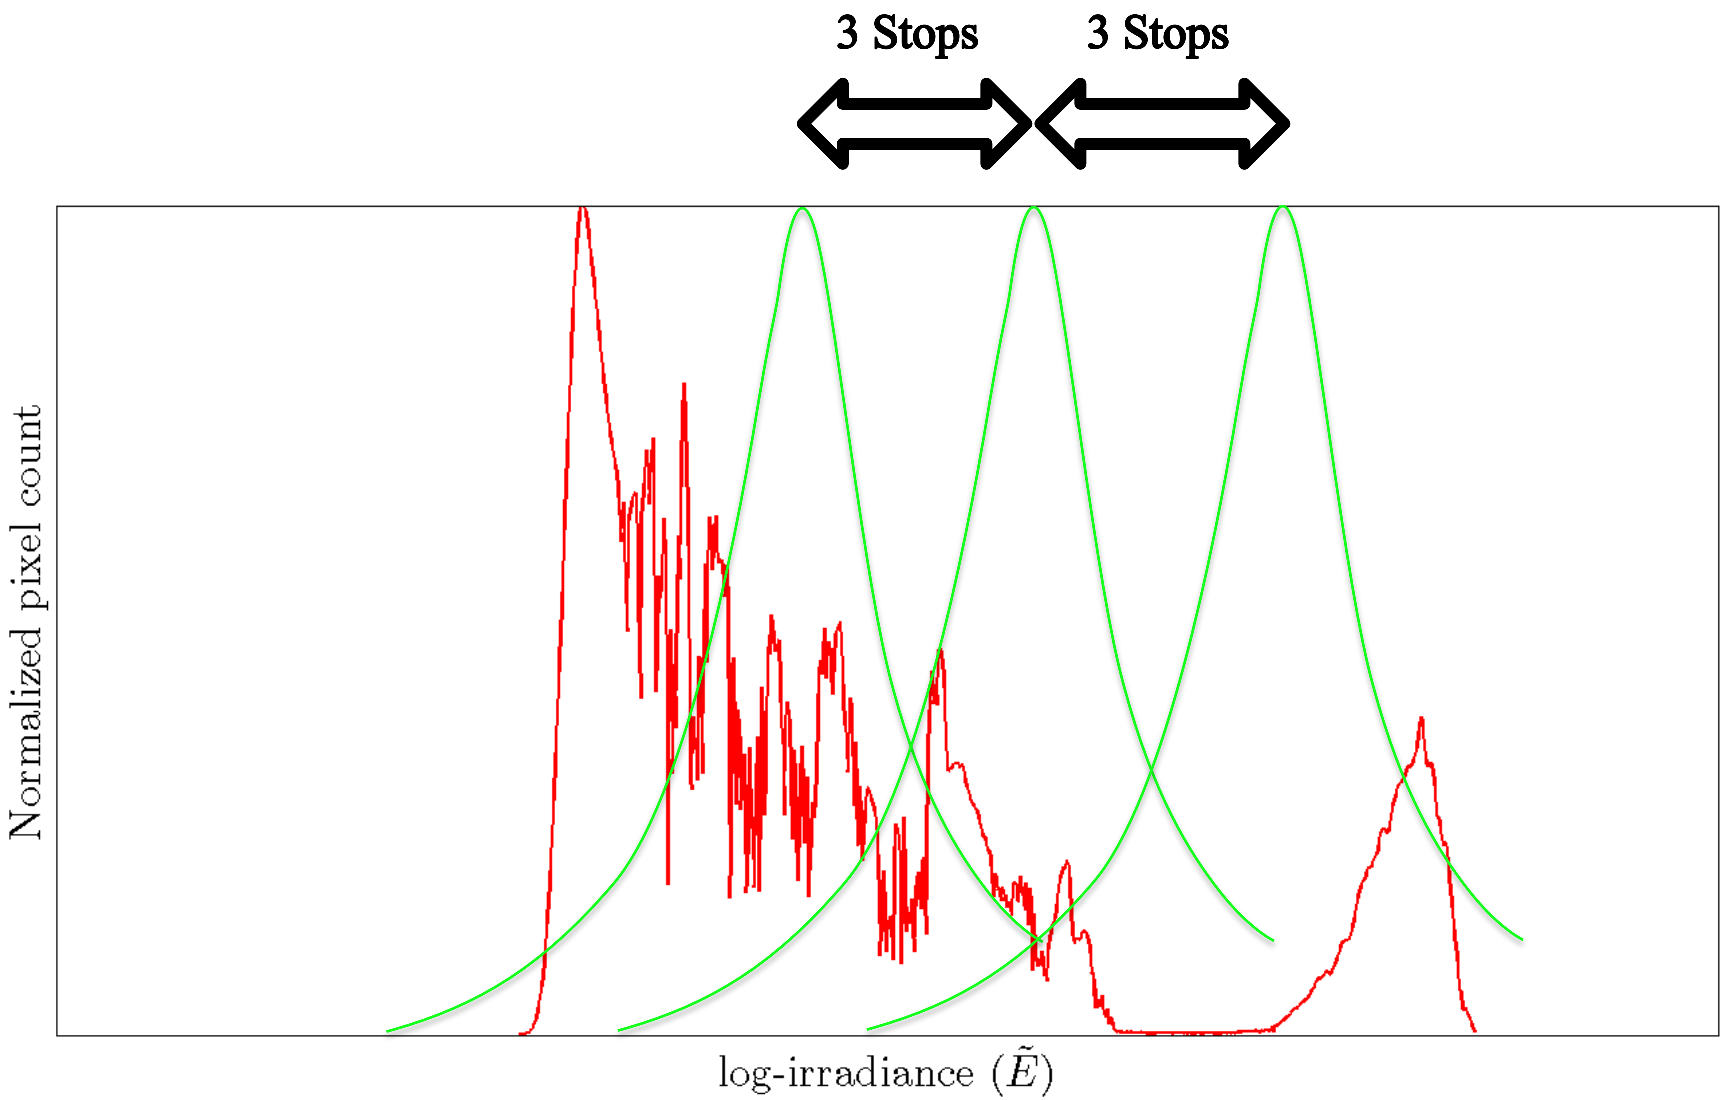
3 Stops
3 Stops
Normalized pixel count
log-irradiance ($\tilde{E}$)

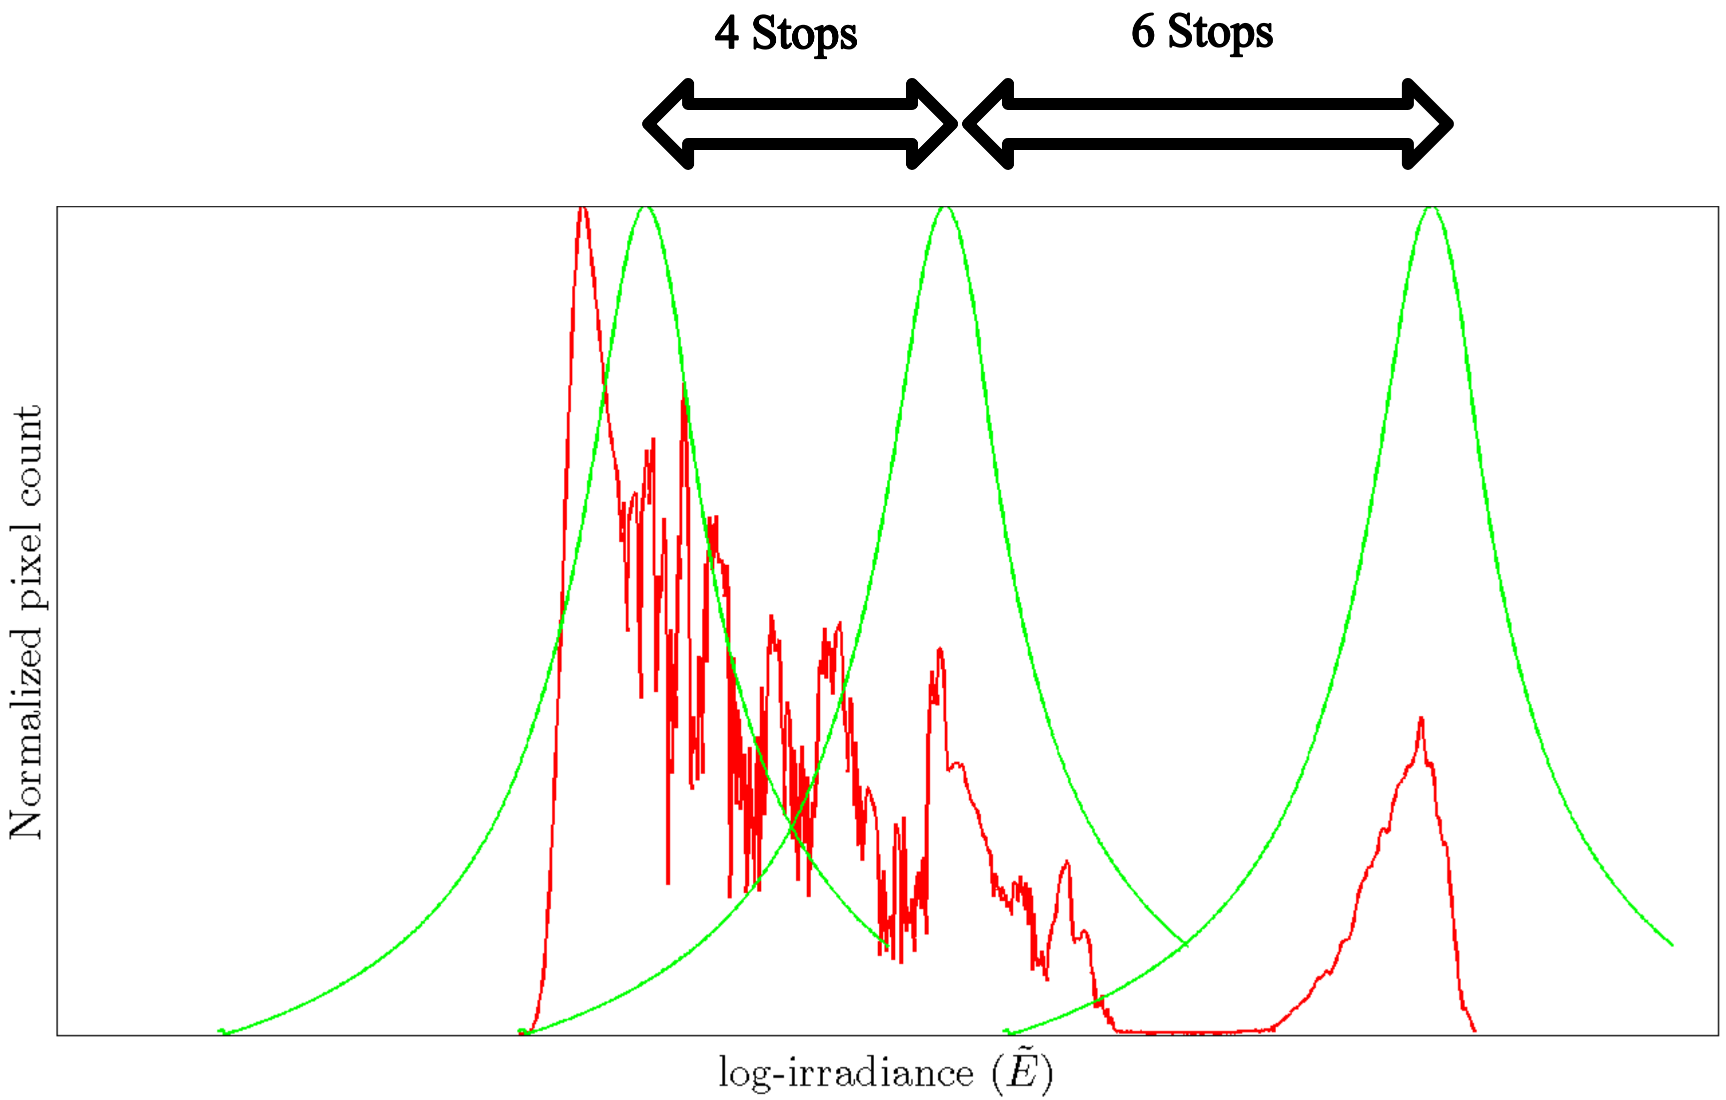
4 Stops
6 Stops
Normalized pixel count
log-irradiance ($\tilde{E}$)

NVIDIA Tegra 3 implementation

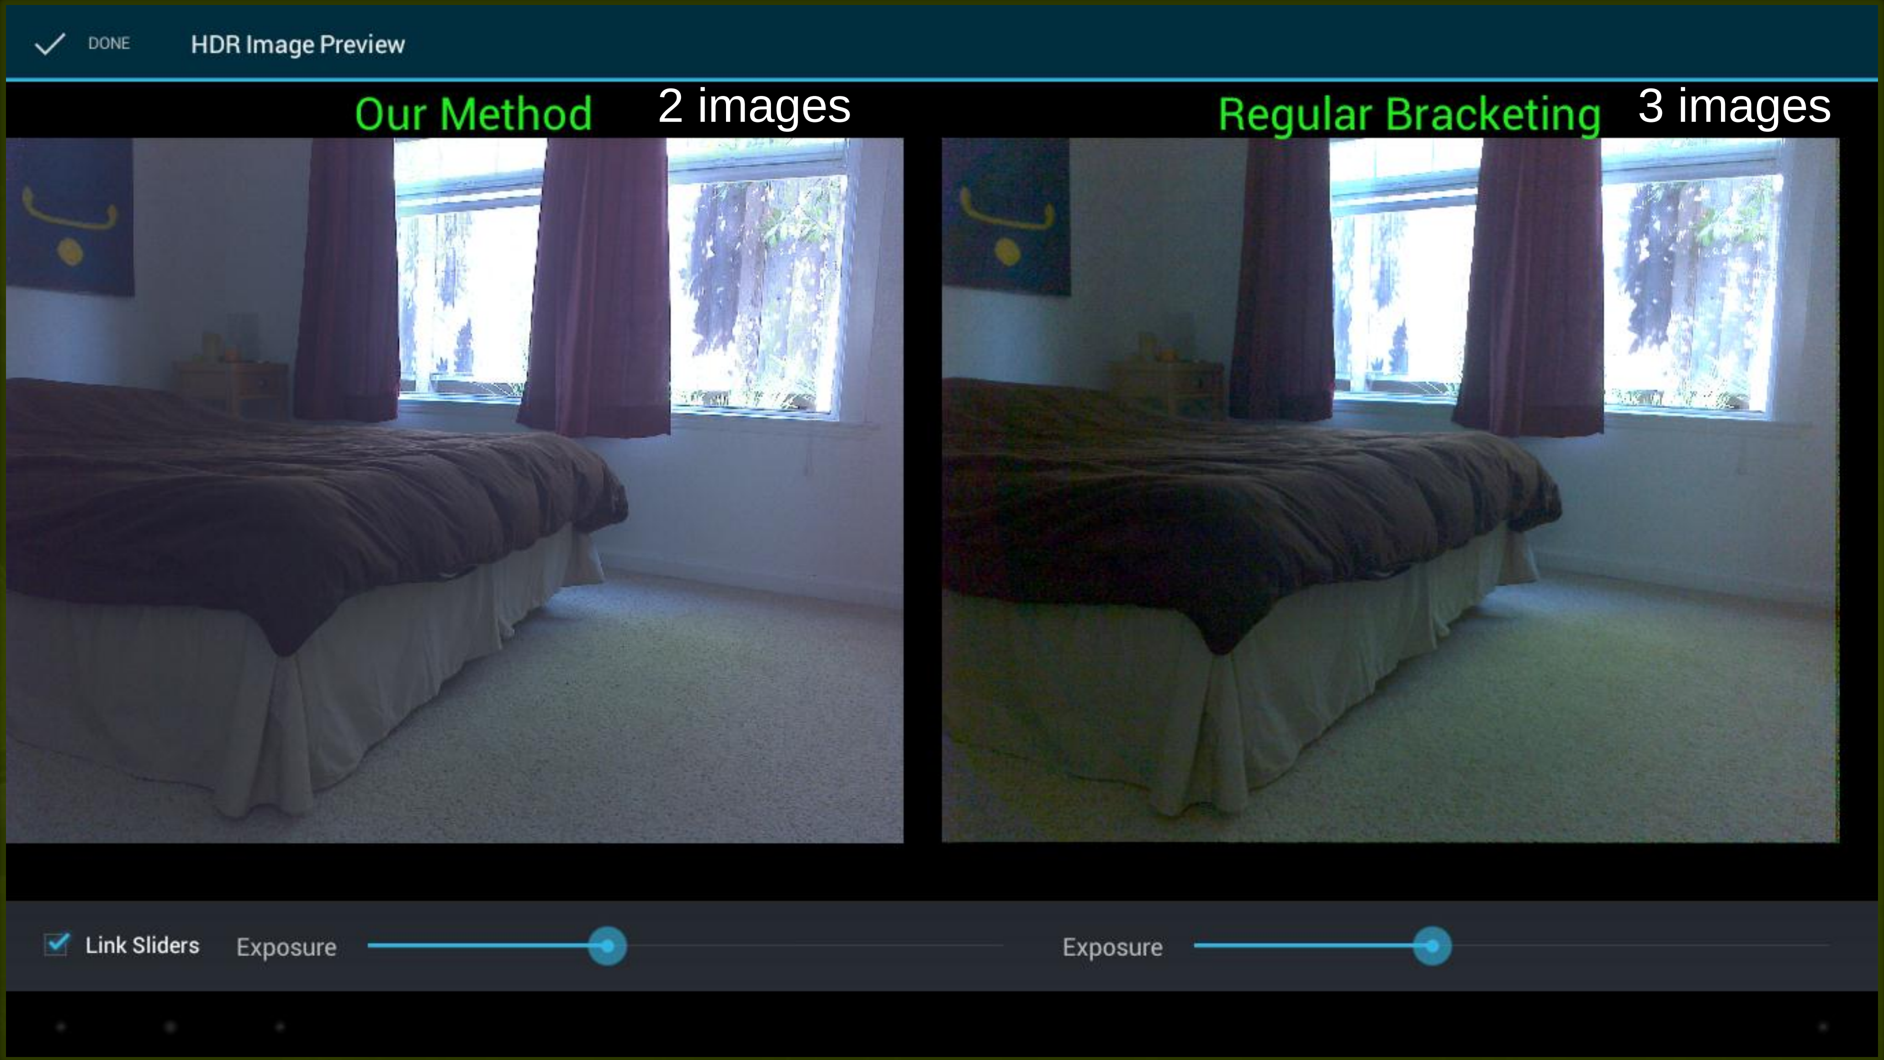

NVIDIA Tegra 3 implementation

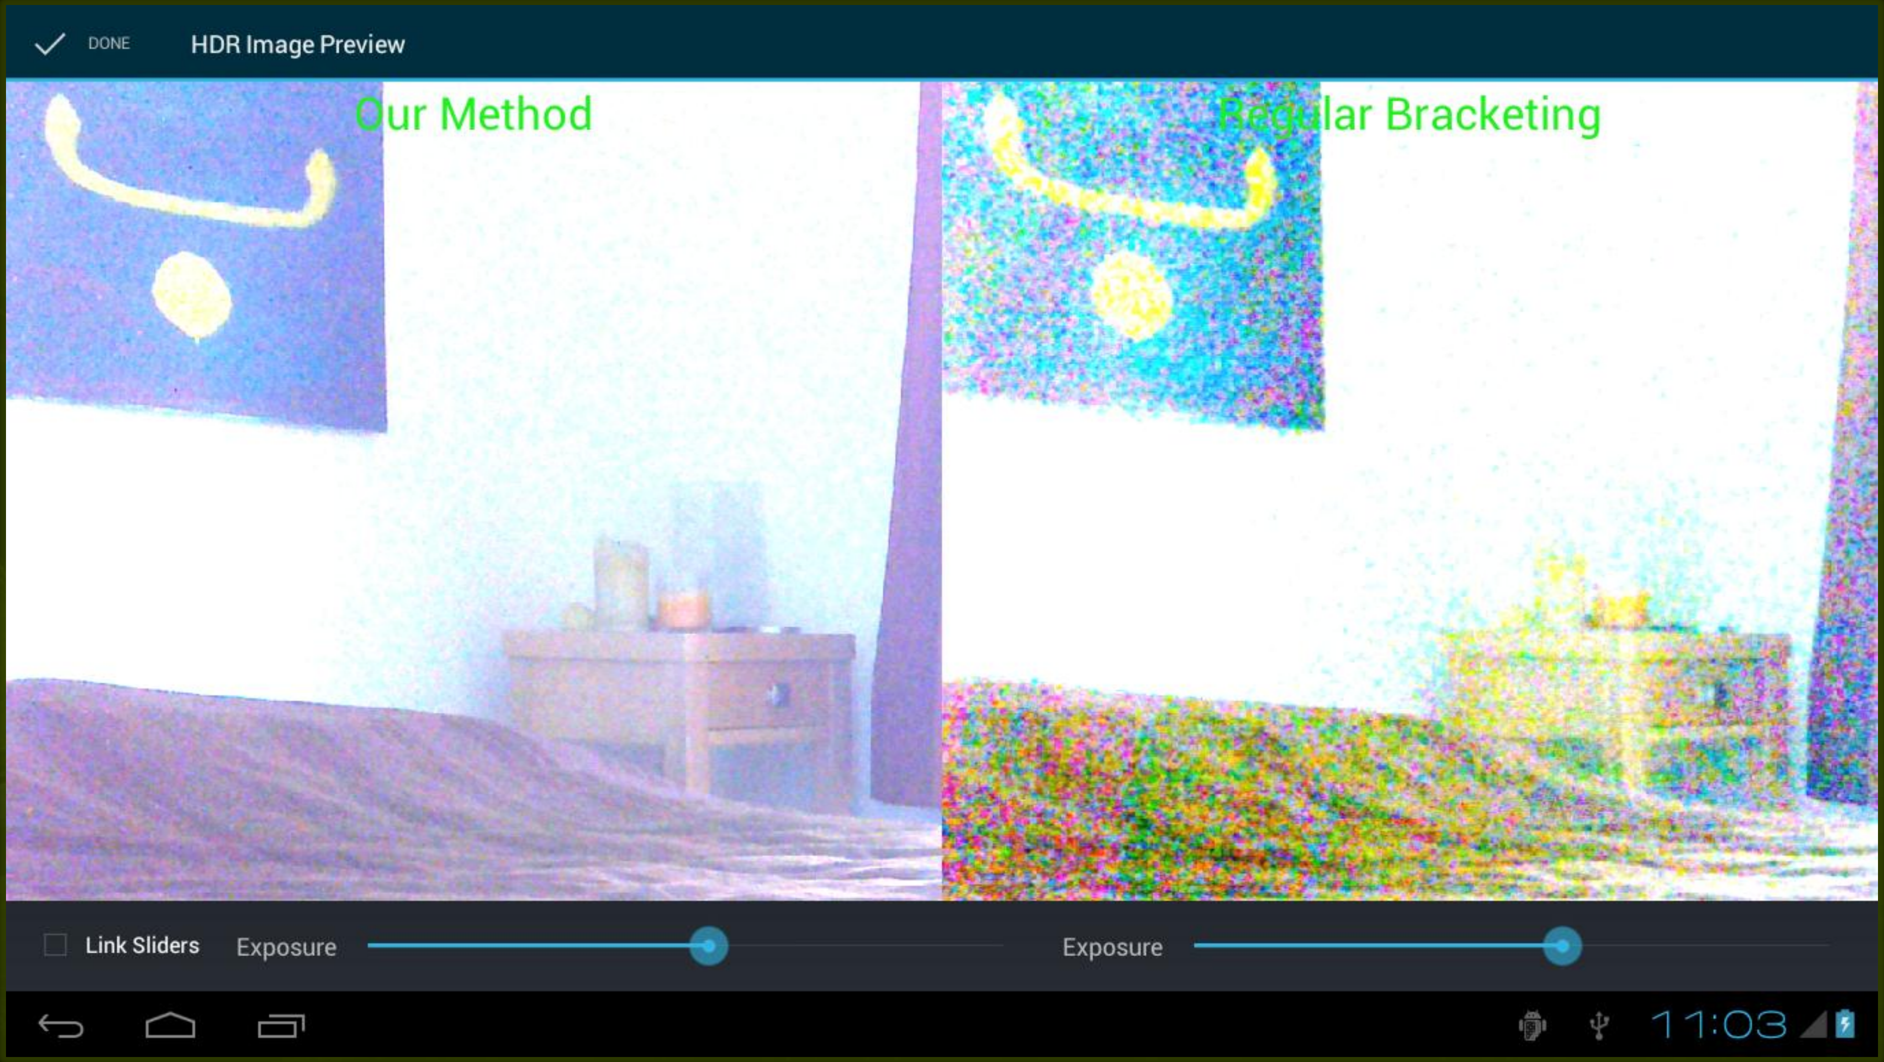

NVIDIA Tegra 3 implementation

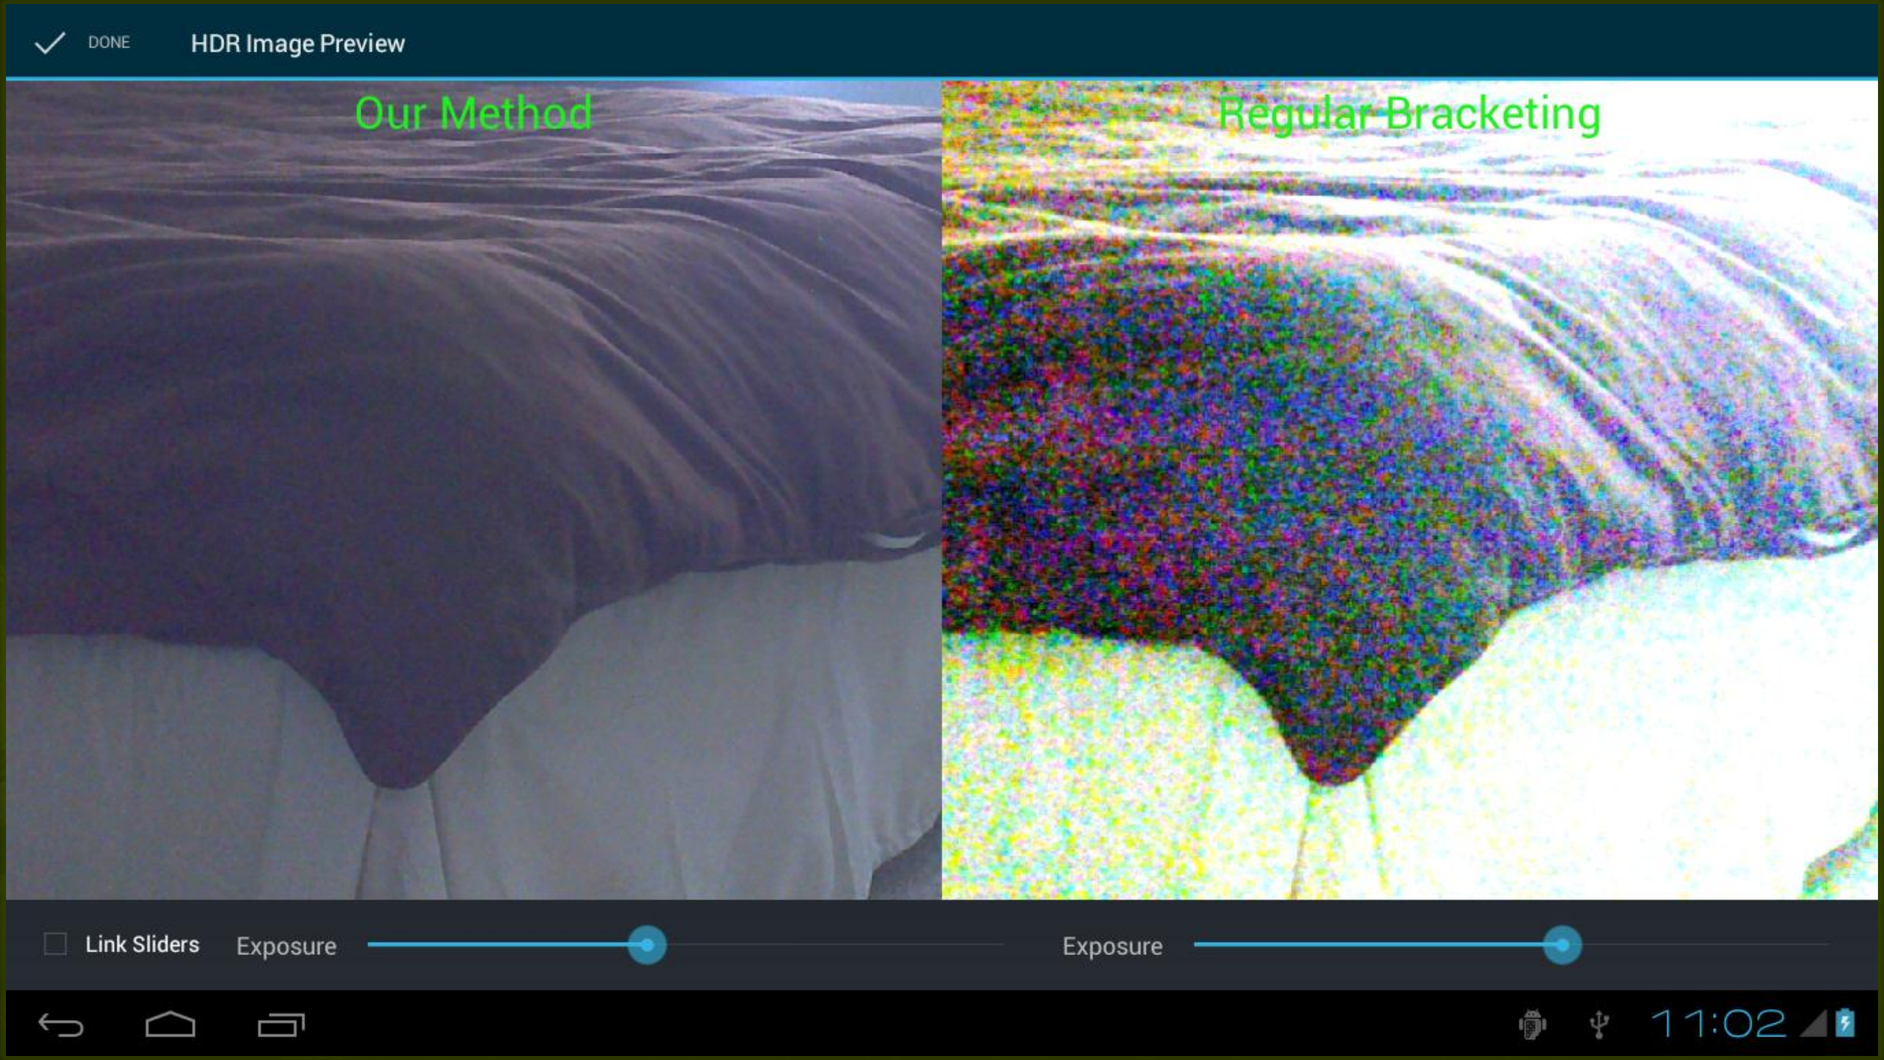

- Fast linear-time performance
  - For 5MP images
  - Global TMO: 5ms
  - Local TMO: 80ms
  - Exposure fusion: ~200ms

graphics.stanford.edu/courses/cs478/

CS 478 - Computational photography

Winter, 2012

A cutaway view showing some of the optical and electronic components in the Canon 5D, a modern single lens reflex (SLR) camera. In the first part of this course, we'll take a trip down the capture and image processing pipelines of a typical digital camera.

This is the Stanford Frankencamera, an experimental open-source camera we are building in our laboratory. It's bigger, heavier, and uglier than the Canon camera, but it runs Linux, and its metering, focusing, demosaicing, denoising, white balancing, and other post-processing algorithms are programmable. We'll eventually be distributing this camera to researchers worldwide.

This is the Nokia N900, the firs in a new generation of Linux-based cell phones. It has a 5-megapixel camera and a focusable Carl Zeiss lens. More importantly, it runs the same software as our Frankencamera, so it's programmable right dow to its autofocusing algorithm.

This is a prototype Nvidia tablet featuring the Tegra 3 processor. It has stereo back-facing cameras, Android OS, and a ported implementation of our FCam API. Each student will receive a tablet for the duration of the course, to try his hands at mobile computational photography.

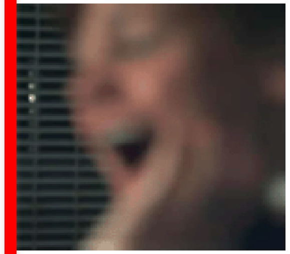

In the second part of the course, we'll consider problems in photography and how they can be solved computationally. One such problem is misfocus. By inserting a microlens array into a camera, one can record light fields. This permits a snapshot to be refocused after capture.

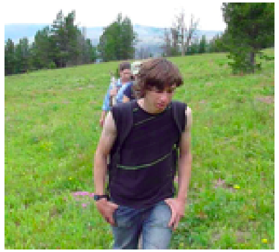

Most digital cameras capture movies as well as stills, but handshake is a big problem, as exemplified by the home video above. Fortunately, stabilization algorithms are getting very good; look at this experimental result. We'll survey the state-of-the-art in this evolving area.

Quarter
Winter, 2012

Units
3-4 (same workload) (+/NC or letter grade)

Time
Mon/Wed 2:30 - 3:45

Place
392 Gates Hall (graphics lab conference room)

Course URL
cs478.stanford.edu

Discussion
CS478 @Piazza

Instructors
Jongmin Baek, Dave Jacobs, Kari Pulli (Guest Lecturer)

Office hours
Wed 3:45 - 5:00, Thurs 2:30 - 3:45, Gates 360

Prerequisite
An introductory course in graphics or vision, or CS 178; good programming skills

Televised?
No

Non-photorealistic viewfinder

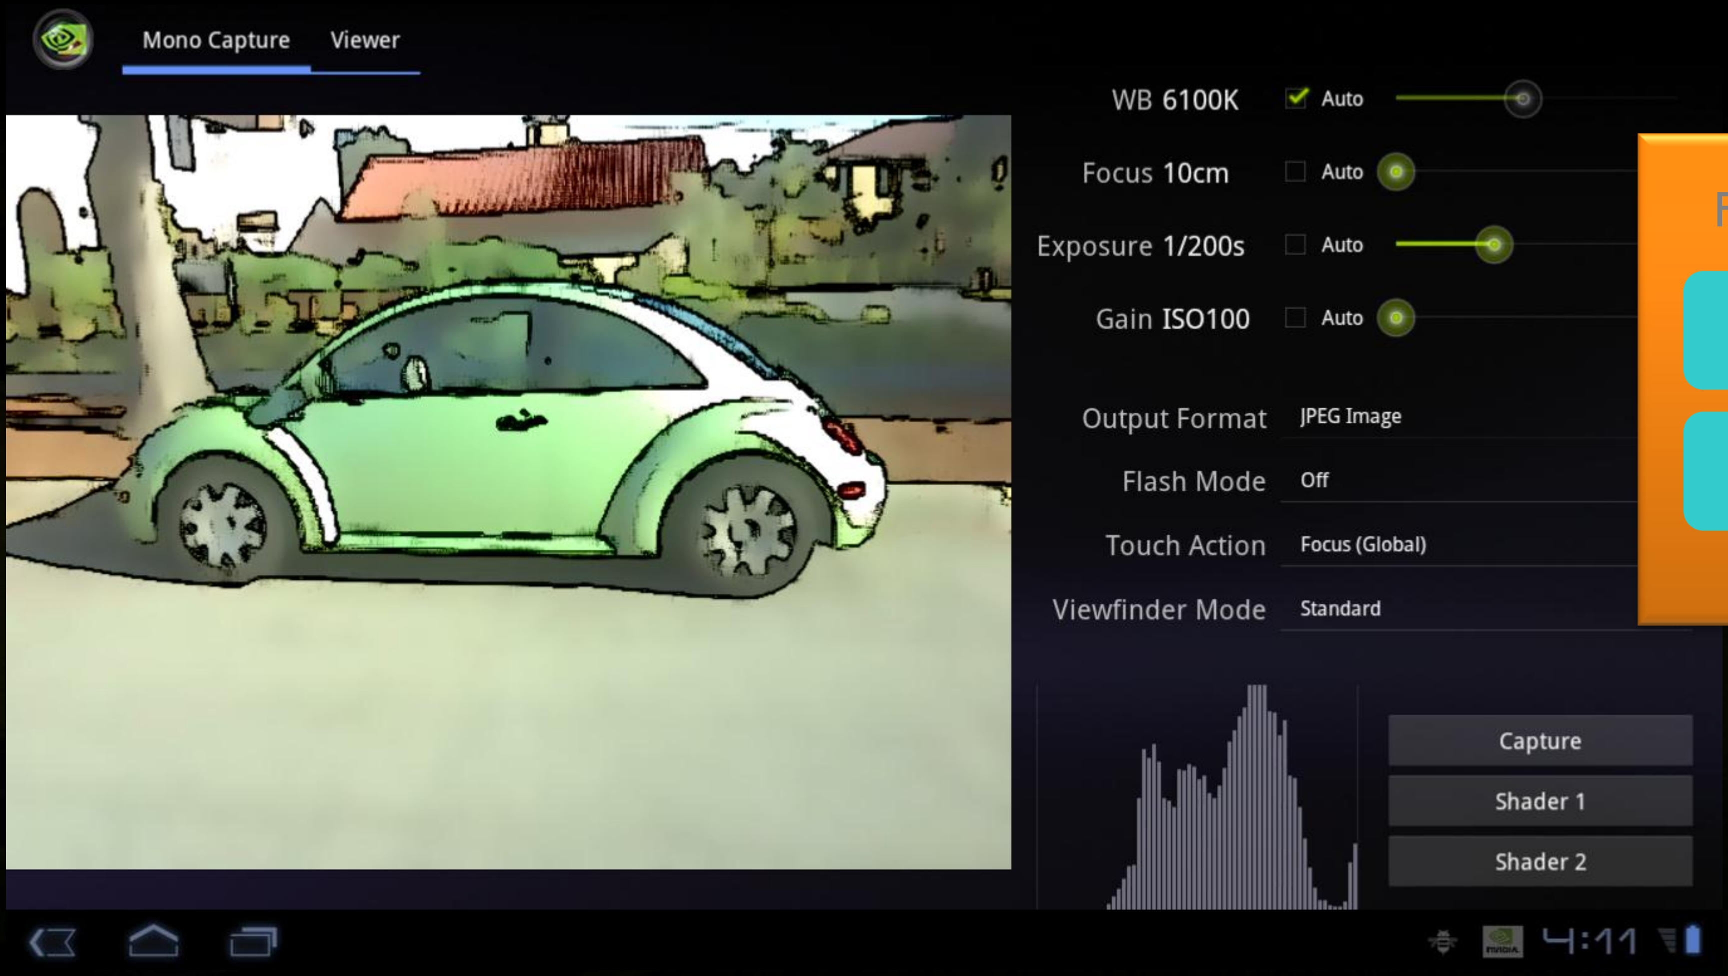

Fragment shaders

1. Bilateral Filter
2. Edge detection

CARMA DevKit

CUDA for ARM Development Kit

CUDA GPU Tegra ARM CPU

Tegra 3 Quad-core ARM A9
Quadro 1000M (96 CUDA cores)
Ubuntu

Gigabit Ethernet
SATA Connector
HDMI, USB

Launch in early fall 2012

Pre-register on www.nvidia.com/CARMADevKit

500 Researchers

Around The Globe Interested In Doing Research Using CARMA

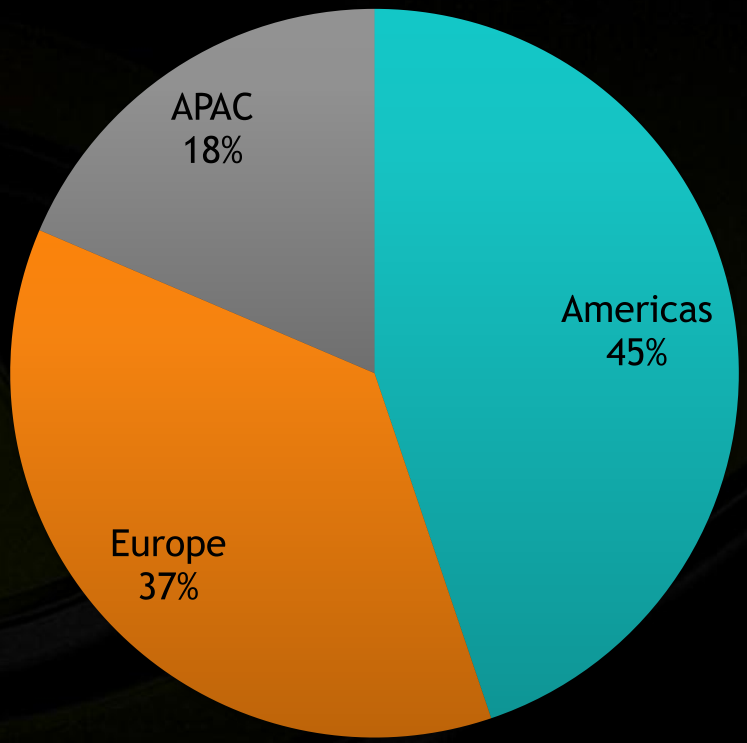

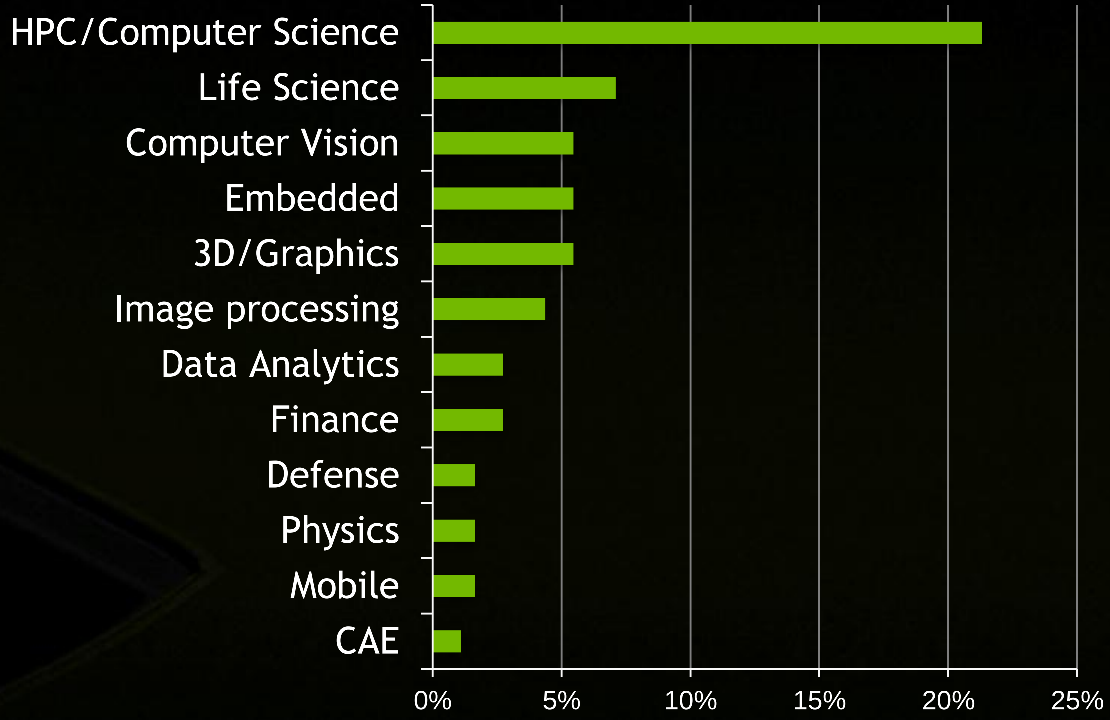

Summary

- Tegra – a great platform for mobile visual computing
- OpenCV – standard CV API optimized for Tegra
- FCam – flexible camera control on Tegra dev board
- CUDA – powerful image processing with Carma

GTC 2013 | March 18-21 | San Jose, CA

The Smartest People. The Best Ideas. The Biggest Opportunities.

Opportunities for Participation:

SPEAK - Showcase your work among the elite of graphics computing

- Call for Sessions: August 2012
- Call for Posters: October 2012

REGISTER - learn from the experts and network with your peers

- Use promo code **GM10SIGG** for a 10% discount

SPONSOR - Reach influential IT decision-makers

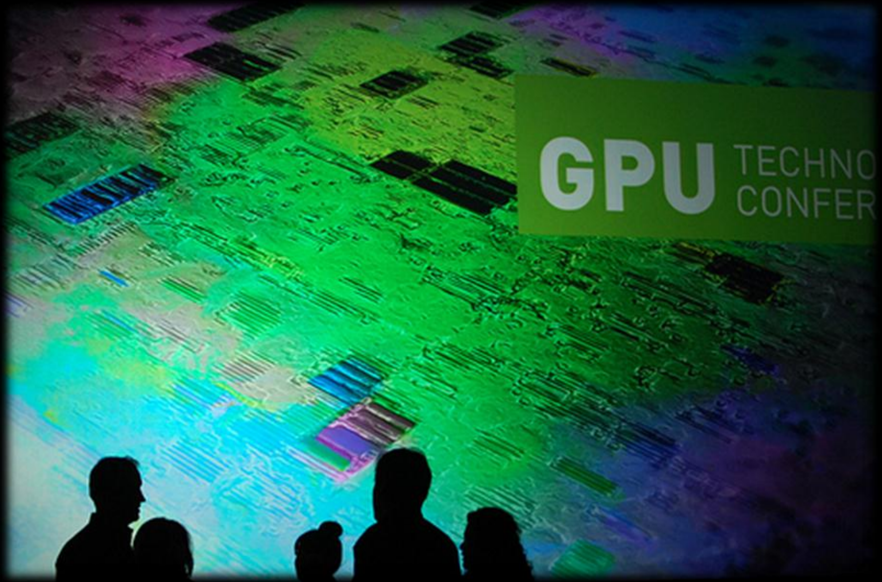

Learn more at www.gputechconf.com