





# TEGRA X1: NVIDIA'S NEW MOBILE SUPER-CHIP

Steve Molnar



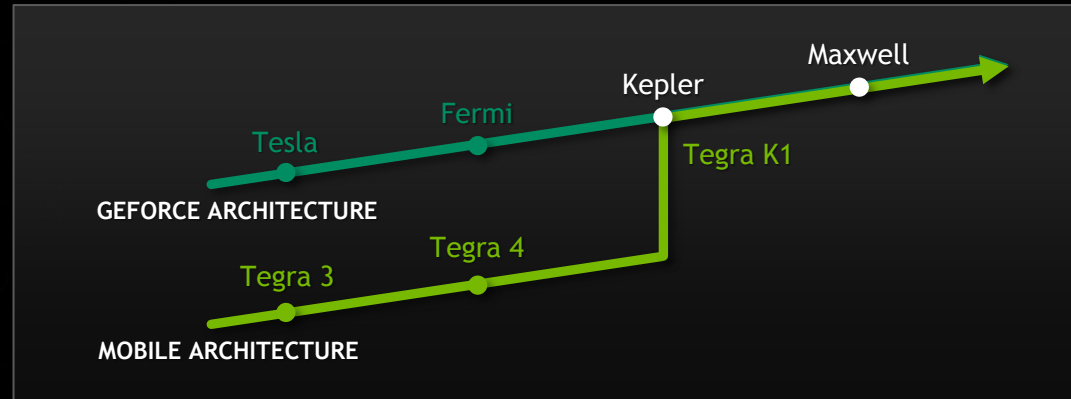
HPG2015 Hot3D




# 2014: TEGRA K1


NVIDIA'S FIRST SUPERCHIP:  
*Console features and performance on a mobile processor*

- *Kepler GPU*
- *192 CUDA cores*



# GPU TRAJECTORIES: DISCRETE AND MOBILE



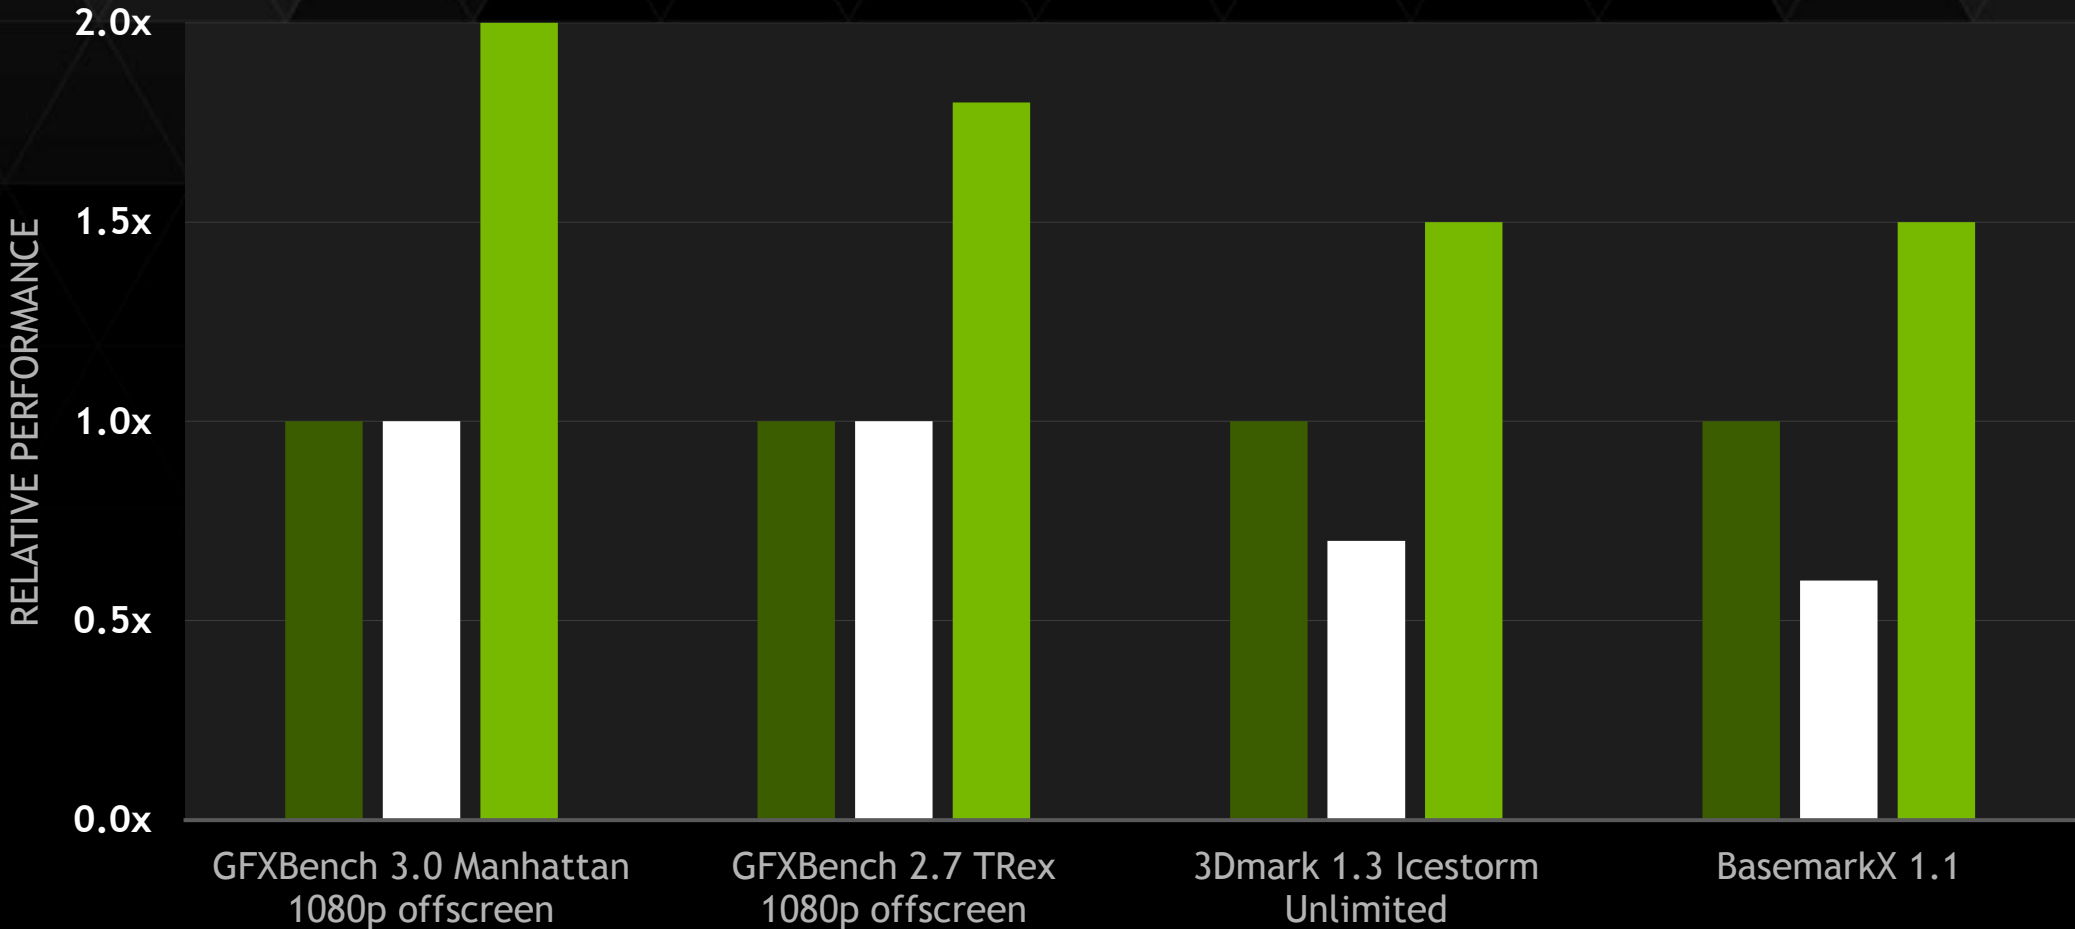


# TEGRA X1: NVIDIA'S NEW MOBILE SUPERCHIP

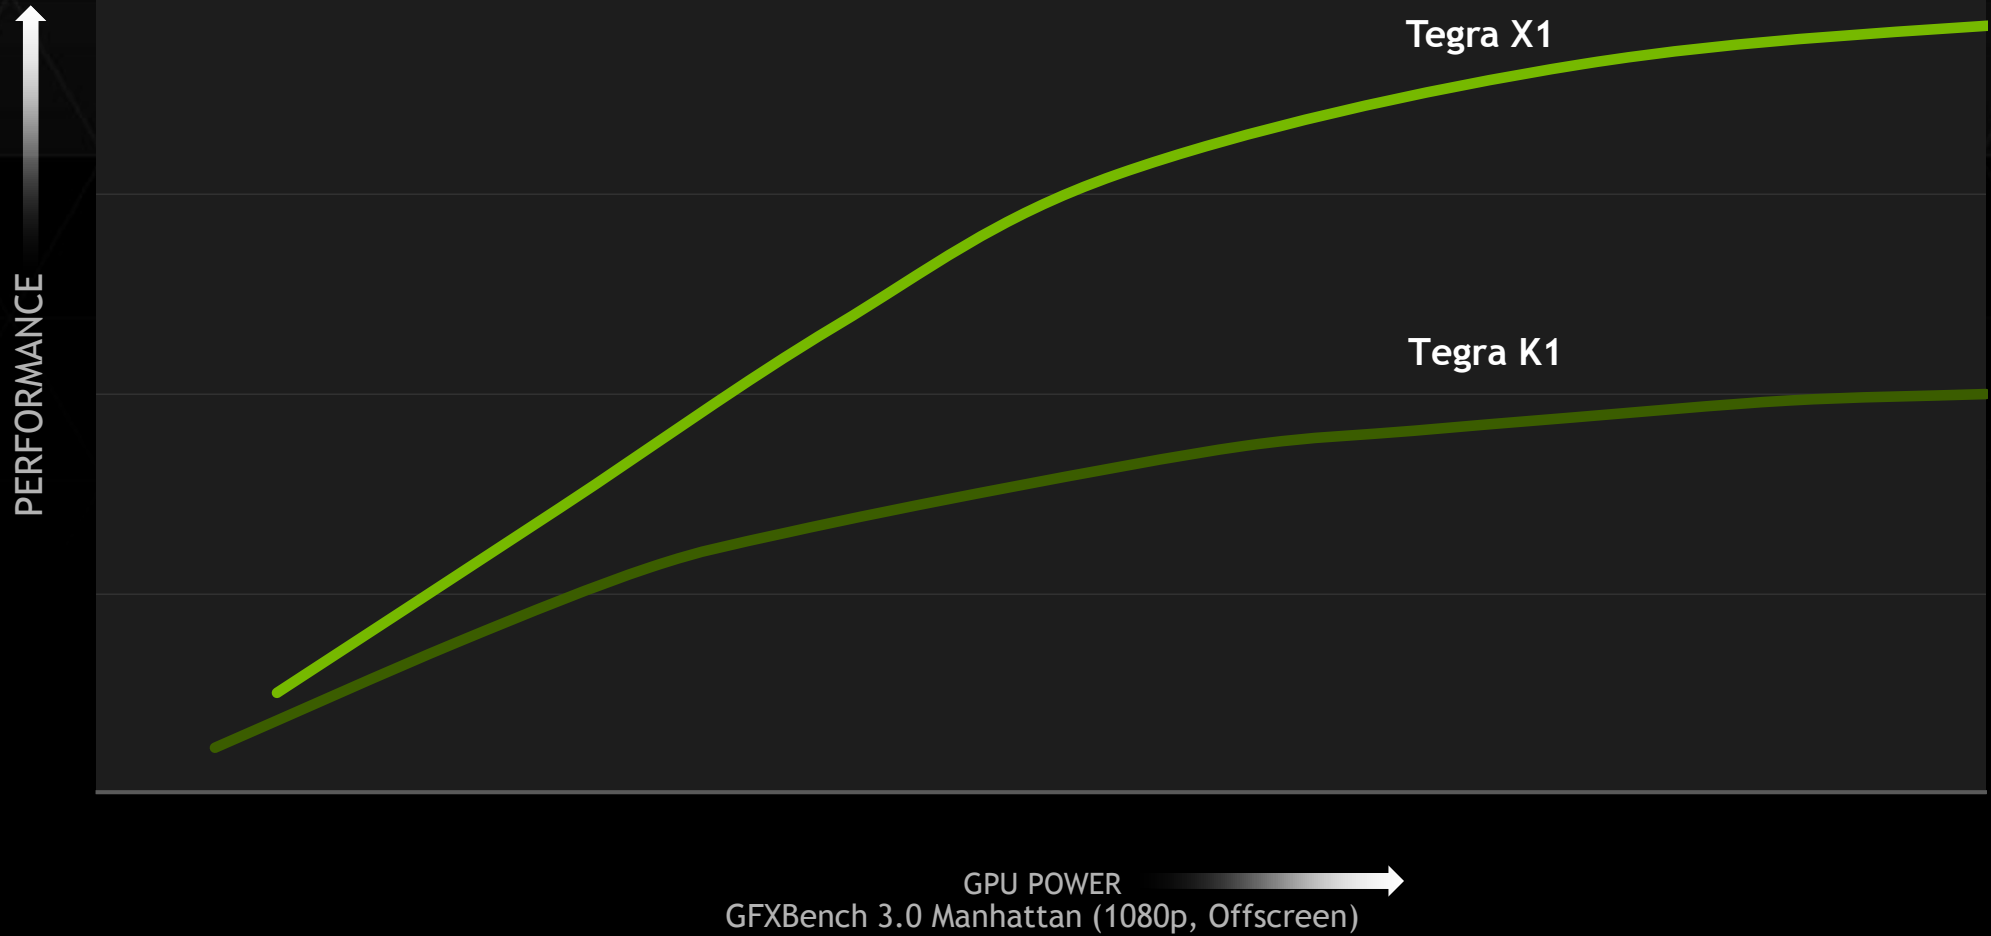
256-core Maxwell GPU | 8-core 64bit ARM CPU | 60 fps 4K Video (H.265/VP9) | 20nm

# RAISING THE BAR. AGAIN.

Tegra K1 A8x Tegra X1




# 2X ENERGY EFFICIENT



# WORLD'S 1<sup>ST</sup> TERAFLOPS MOBILE PROCESSOR








# THE MOST ADVANCED GPU. AGAIN.

	TEGRA X1 Maxwell Graphics	GEFORCE GTX 980
OpenGL ES 3.1	✓	✓
AEP	✓	✓
OpenGL 4.5	✓	✓
DX12	✓	✓
CUDA 6.0	✓	✓



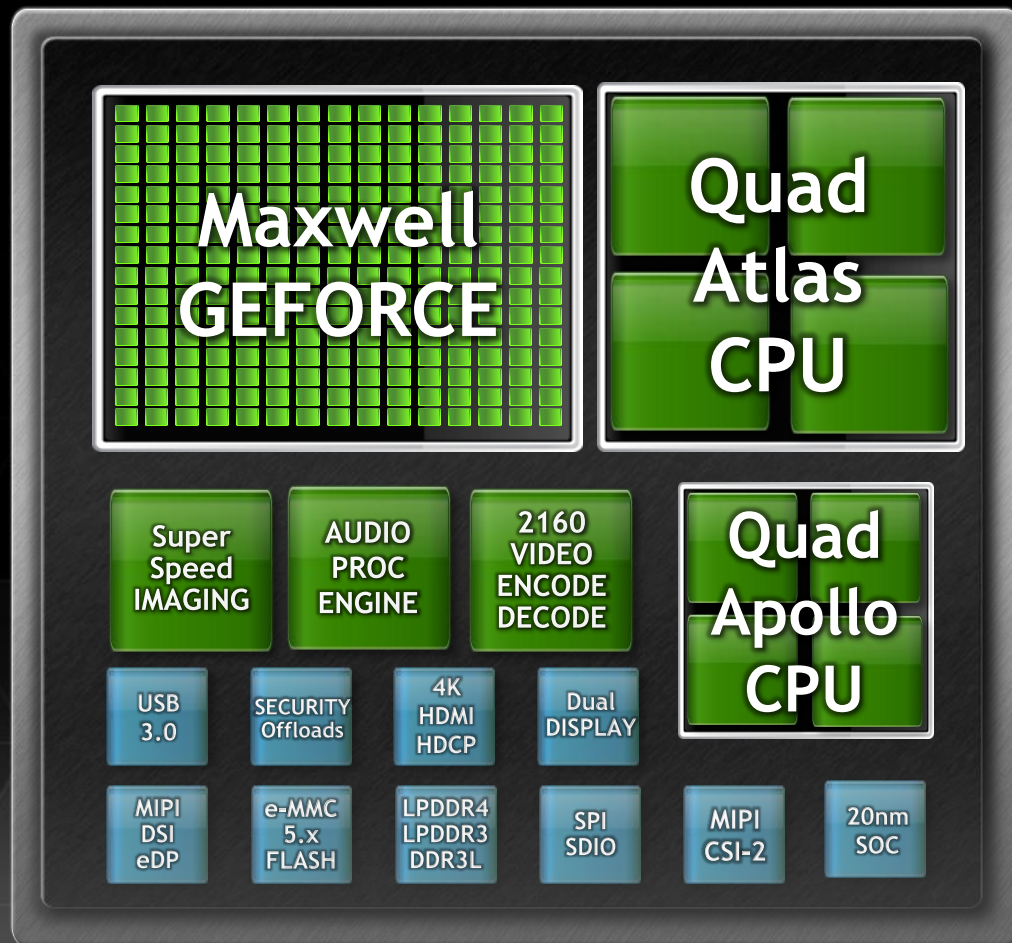
REST  
OF THIS  
TALK

- 1 Processor and memory architecture
- 2 Maxwell GPU advances
- 3 Computer vision
- 4 4K video
- 5 Shield Android TV / console
- 6 Demos



# TEGRA X1 PROCESSOR AND MEMORY ARCHITECTURE

# TEGRA X1 - OVERVIEW



GRAPHICS	Maxwell GeForce - World's Fastest GPU <i>2 x SMM units, DX-next, OpenGL 4.4</i>
CPU	Octo-Core 64b ARM v8 CPU Complex <i>4xCA57 Atlas/2MB L2; 4xCA53 Apollo/512KB L2</i>
MEMORY	64b / Dual-Quad Channel Memory Interface <i>LPDDR4-3200, LPDDR3E-1866, DDR3L-1866</i>
VIDEO	4K x 2K Encode and Decode <i>H.264, H.265, VP8, VP9 (dec-only)</i>
POWER	Low Power <i>20nm, HW Offloads, Isolated Pwr Rails, PRISM</i>
DISPLAY	4K x 2K 24b @60Hz, 1080p @120Hz <i>DSI 2x4, eDP, High Speed HDMI 2.0, DP</i>
IMAGING	Full Quad Camera imaging, Dual ISP 650Mp/s <i>Maxwell 16fp imaging GPGPU, HW AO-HDR</i>
Mobile I/O	Designed for mobile <i>e.MMC5.x, USB3.0/2.0/HSIC, SD/SDIO 3.0, CSI-2</i>


# TEGRA X1 CPU CONFIGURATION

## 4 HIGH PERFORMANCE A57 BIG CORES


- ▶ 2MB L2 cache
- ▶ 48KB L1 instruction cache
- ▶ 32KB L1 data cache

## 4 HIGH EFFICIENCY A53 LITTLE CORES

- ▶ 512KB L2 cache
- ▶ 32KB L1 instruction cache
- ▶ 32KB L1 data cache



# 2X CPU EFFICIENCY



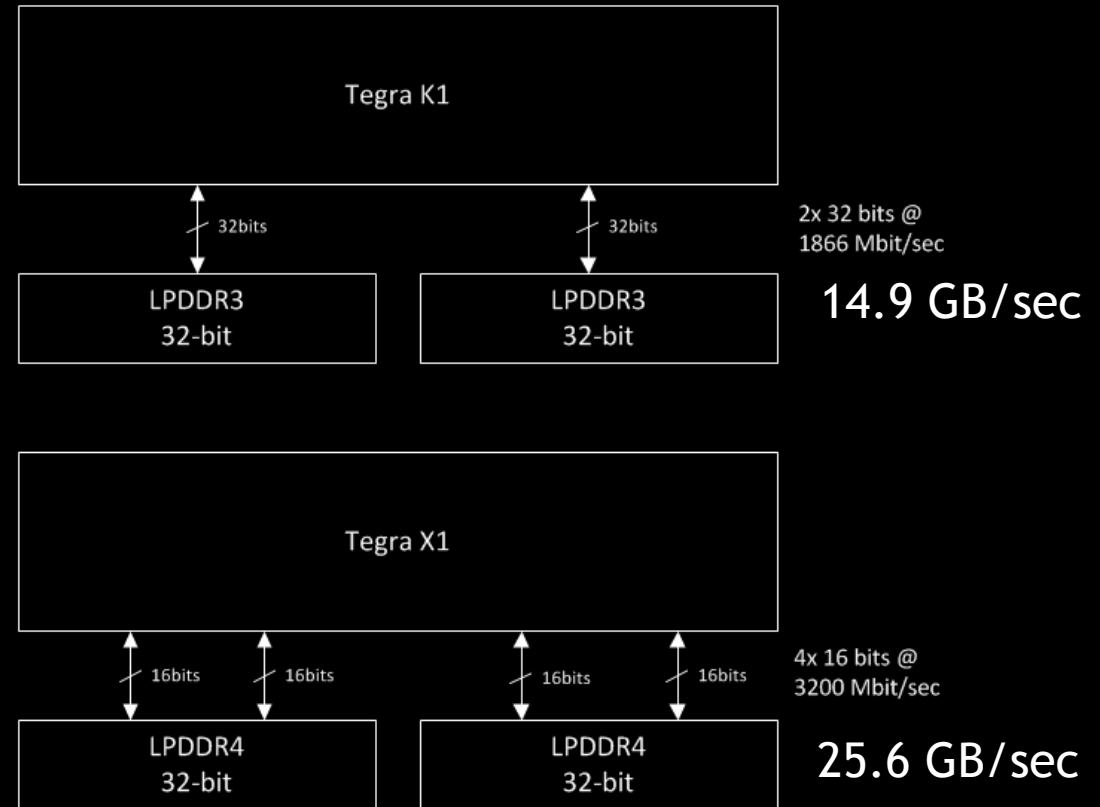
SPECint score estimates. Both Tegra X1 and Exynos 5433 have ARM A57 cores, and are manufactured using 20nm process. Measured estimates is using SPEC CPU2000. Tegra X1 is measured on NVIDIA reference platform and Exynos 5433 is measured on Samsung Note 4

# LPDDR4

- ▶ Memory system capabilities dictated by DRAM technology
- ▶ Internal DRAM array structure and speeds change little over generations (~200 MHz internal clock)
- ▶ IO speed increasing exponentially over last generations
  - ▶ LPDDR 400 Mbit/sec
  - ▶ LPDDR2 1066 Mbit/sec
  - ▶ LPDDR3 1866 Mbit/sec
  - ▶ LPDDR4 3200 Mbit/sec
- ▶ LPDDR4 runs at 40% less power than LPDDR3

# LPDDR4


- ▶ 32-bit channel becomes pair of 16-bit channels with higher burst length (same 32B transfer atom)
- ▶ Low Voltage Swing Terminated Logic (LVSTL) I/O interface
- ▶ Doubling of the interface speed requires additional training, clocking changes



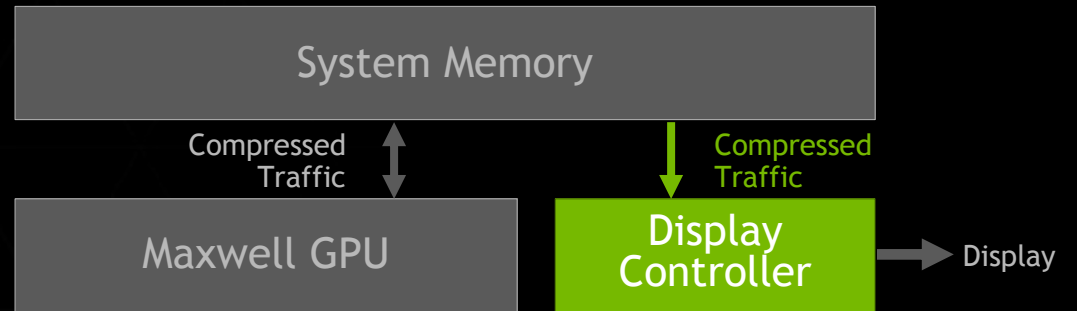


# TEGRA X1 COMPRESSION ADVANCES

## 3<sup>RD</sup> GEN DELTA COLOR COMPRESSION



## NEW END-TO-END COMPRESSION



# MEMORY COMPRESSION

Original Image



*Image taken on Half Life 2*

# MEMORY COMPRESSION

3<sup>rd</sup> Generation  
Delta Color Compression

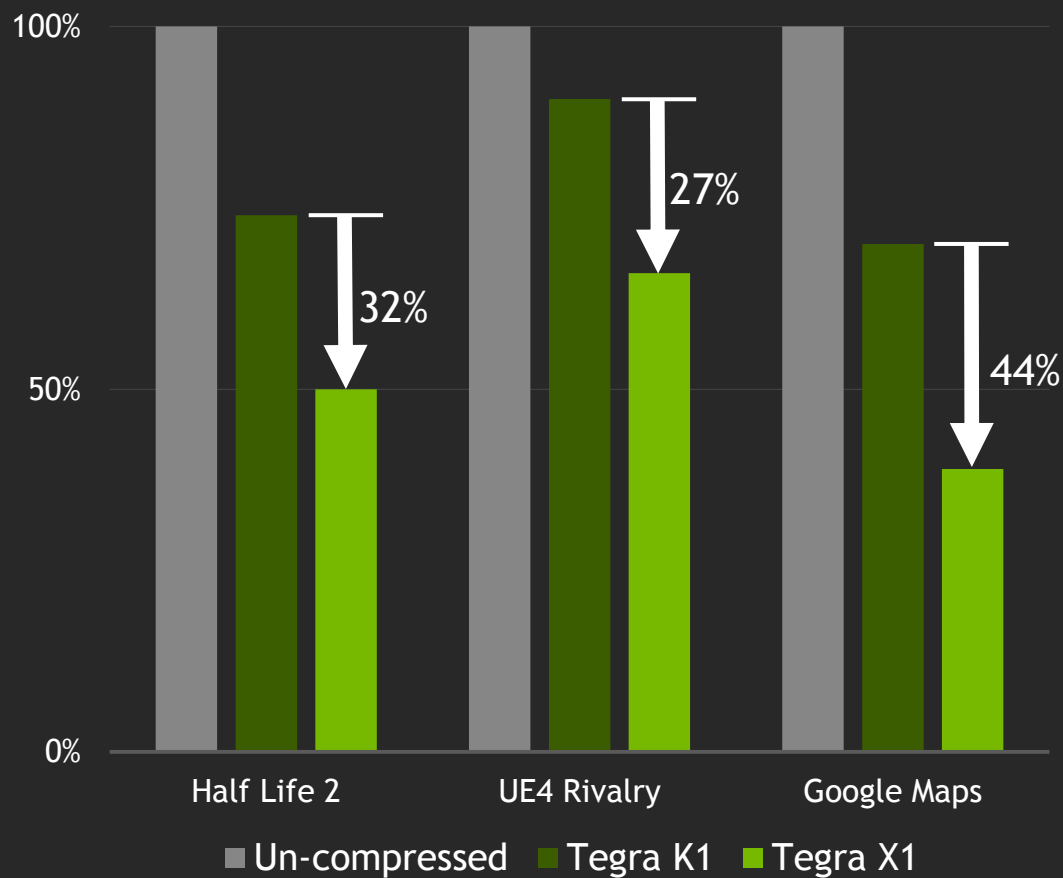


*Compressed pixels colored in purple*

# MEMORY BANDWIDTH REDUCTION

- ▶ Enhanced Compression Algorithms
- ▶ Enhanced Caching Effectiveness
- ▶ New End-to-End Compression

## Memory Bandwidth Reduction



# Tegra X1 DRAM Energy Efficiency

GFXBench 3.0 Manhattan


	TEGRA K1	TEGRA X1	
Bytes per frame efficiency	1x	1.5x	Bandwidth reduction due to Maxwell architecture + end-to-end compression
Dram pJ/byte efficiency	1x	1.4x	pJ/Byte reduction due to LPDDR4 vs LPDDR3
DRAM energy efficiency	1x	2.1x	



TEGRA X1  
MAXWELL GPU

# MAXWELL

THE MOST ADVANCED GPU EVER BUILT




*“...faster, less power hungry, and quieter...  
once again NVIDIA has landed the technical trifecta”*

— Anandtech

# TEGRA X1 MAXWELL GPU

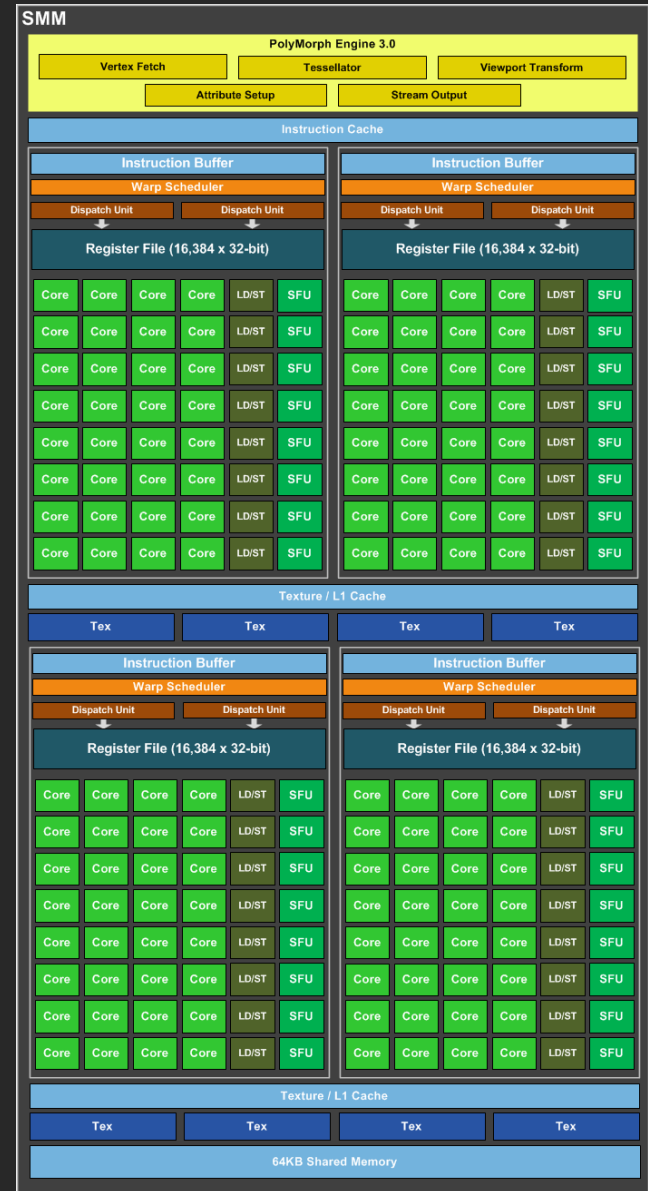
- ▶ 2x performance vs Tegra K1
- ▶ 2x perf/watt vs Tegra K1
- ▶ 2 SM
- ▶ 256 CUDA Cores
- ▶ 2 Geometry Units
- ▶ 16 Texture Units
- ▶ 16 ROP Units
- ▶ Maxwell Memory Architecture
- ▶ 64-bit LPDDR4





# TEGRA X1 MAXWELL SM

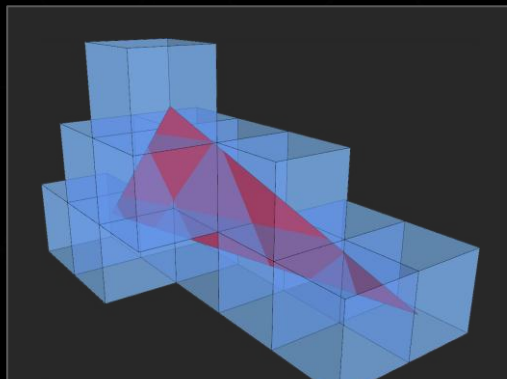
- ▶ 2x Energy Efficiency vs Tegra K1
- ▶ 40% higher perf per core
- ▶ Improved scheduler
- ▶ New datapath organization
- ▶ Enhanced memory hierarchy
- ▶ Double speed FP16




# NEXT-GENERATION GRAPHICS FEATURES




Accelerated Path Rendering




Multi-Projection Acceleration




Raster Ordered View



Conservative Raster



Tiled Resources



Advanced Sampling

# GLOBAL ILLUMINATION

## THE GRAND CHALLENGE OF REAL-TIME COMPUTER GRAPHICS



# VXGI

## VOXEL GLOBAL ILLUMINATION

Dynamic

1-bounce Indirect Diffuse,  
Specular, Reflections, Area Lights

Available for UE4 and  
other major engines (Q4)



# VR DIRECT

Auto Stereo

Low Latency

Auto Asynchronous Warp





# GAMEWORKS™

300 Visual Effects Engineers  
GameWorks Library  
Developer Tools






# TEGRA X1 DEEP LEARNING AND COMPUTER VISION

# BUILT FOR COMPUTER VISION

- ▶ High computational horsepower
- ▶ Low power consumption
- ▶ General purpose programmability
- ▶ Multi-camera support
- ▶ High performance HD video and image processing







- Bicycle
- Street Name
- Pedestrian
- Equipment
- Traffic Cone
- Traffic Cone
- Road Outline
- Hydrant
- Traffic Light
- Van
- Traffic Sign
- Car
- School Bus



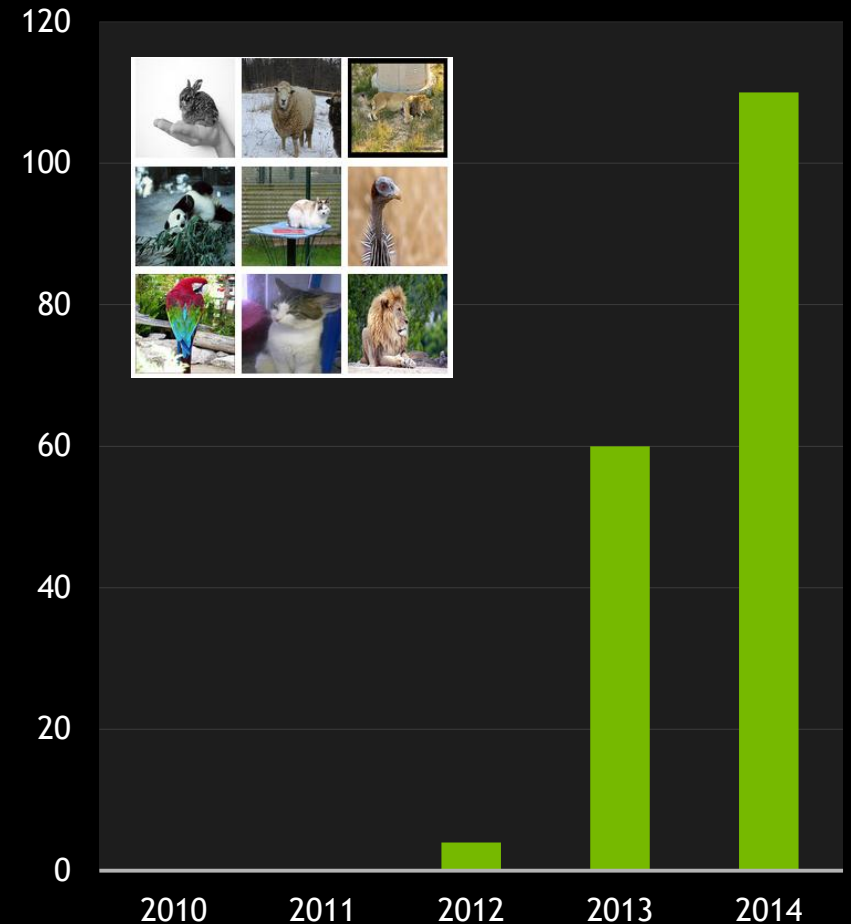
# GPUs Revolutionizing Deep Learning Across Industries

*“In 2009, Andrew Ng and a team at Stanford realized that GPU chips could run neural networks in parallel. Today neural nets running on GPUs are routinely used by cloud-enabled companies such as Facebook to identify your friends in photos or, in the case of Netflix, to make reliable recommendations for its more than 50 million subscribers.”*

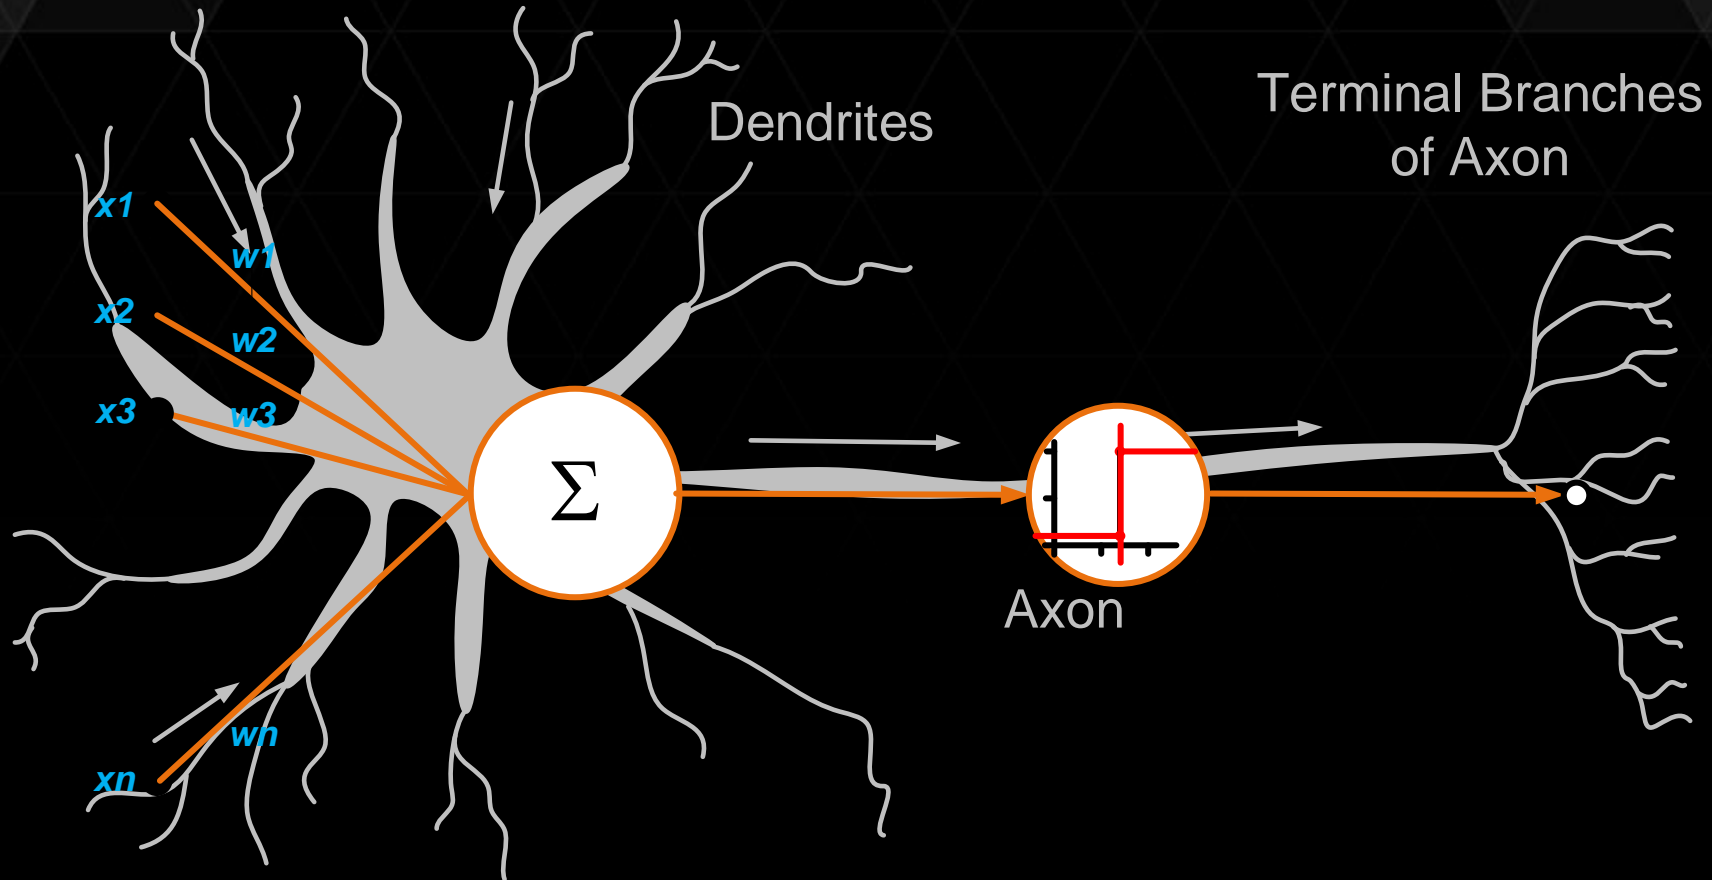
**WIRED**

## 95% OF ENTRIES USING GPUS IN 2014

ImageNet Challenge

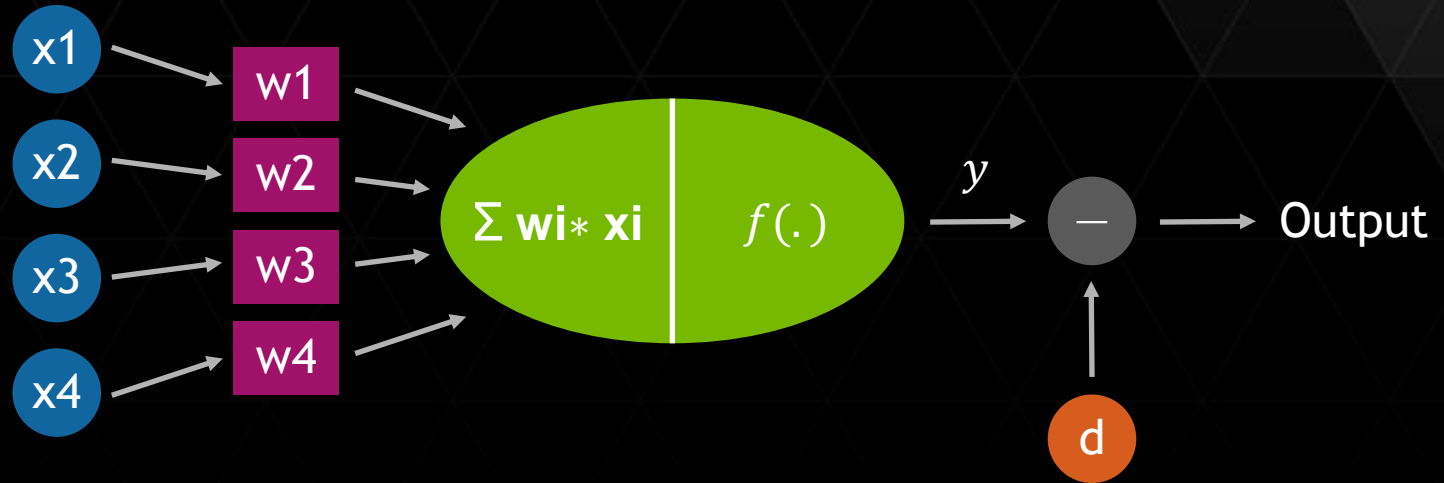


# ARTIFICIAL NEURAL NETWORKS




# THE PERCEPTRON — THE SIMPLEST MODEL

PERCEPTRON




# THE PERCEPTRON – THE SIMPLEST MODEL


## PERCEPTRON




## ACTIVATION FUNCTIONS:



Step Function



Sign Function




Sigmoid Function

## LEARNING:


$$y^{(t)} = f\left\{\sum_i w_i^{(t)} x_i^{(t)}\right\}$$

$$\text{Update} \begin{cases} \Delta w_i^{(t)} = \varepsilon (d^{(t)} - y^{(t)}) x_i^{(t)} \\ w_i^{(t+1)} = w_i^{(t)} + \Delta w_i^{(t)} \end{cases}$$


# MULTILAYER NEURAL NETWORKS



# HOW A TRAINED MACHINE SEES



Image

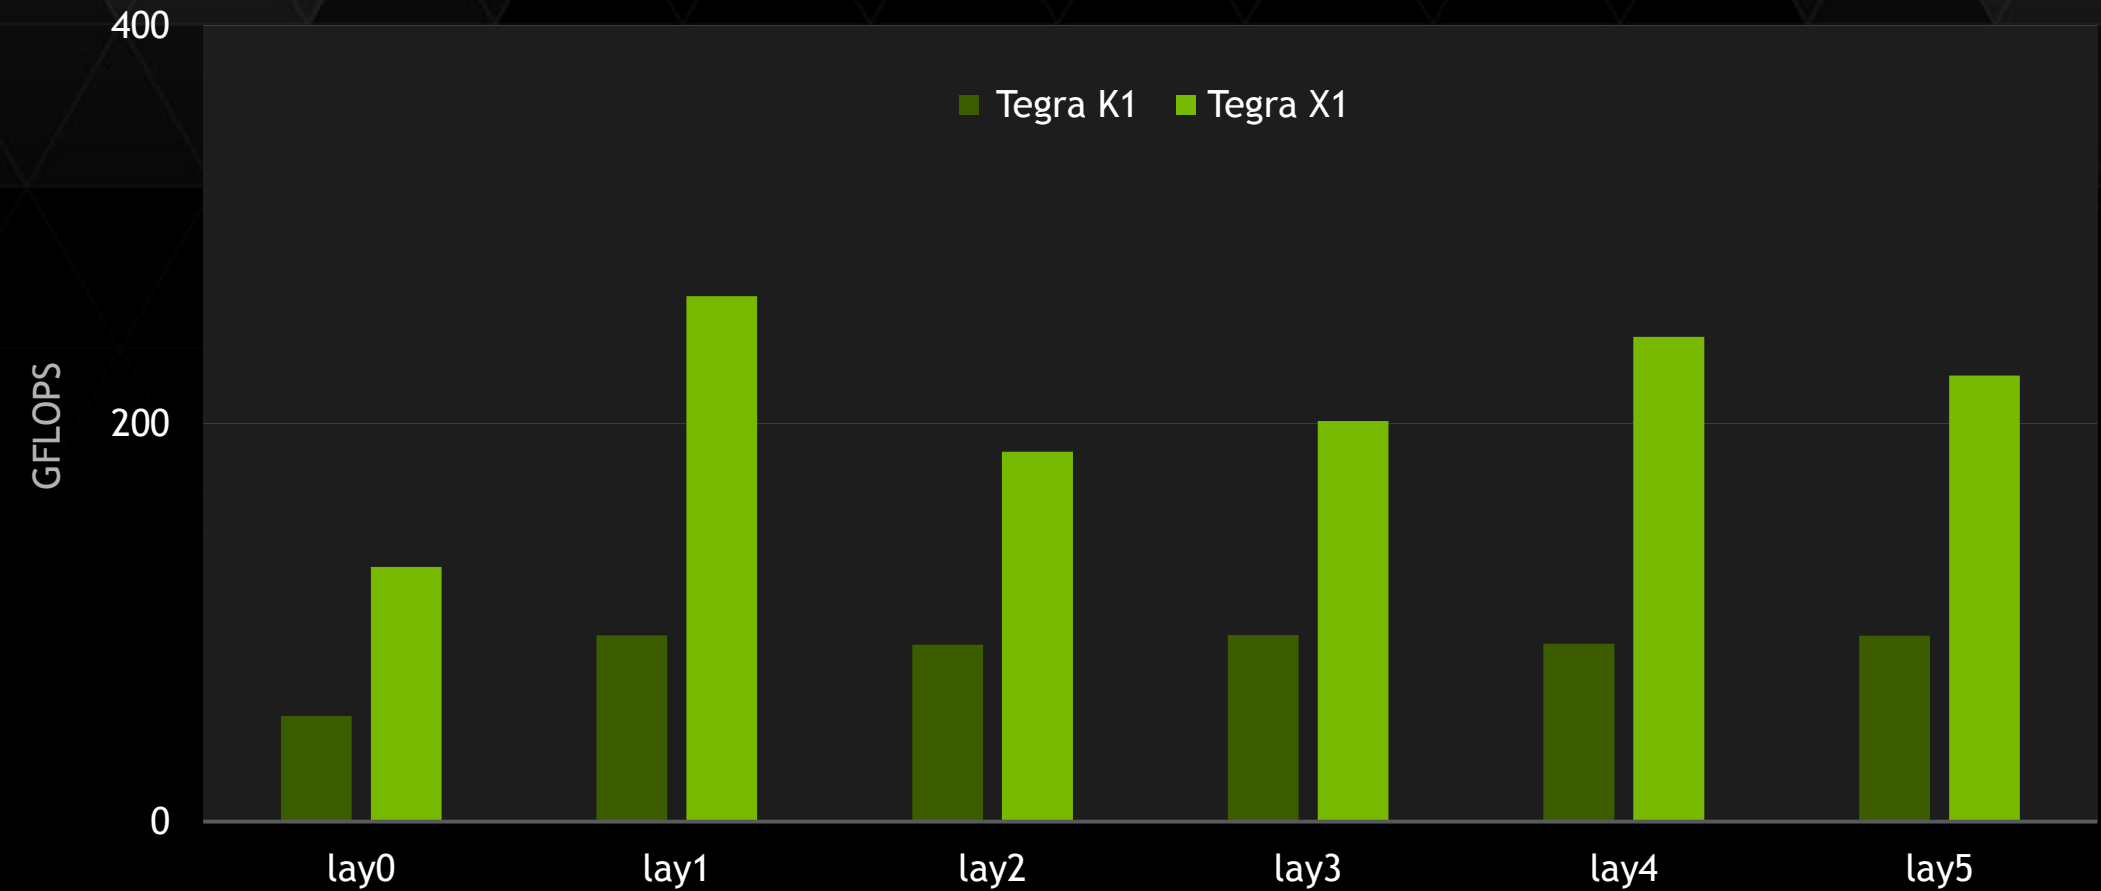


“Audi A7”



# TEGRA X1 FOR DEEP LEARNING

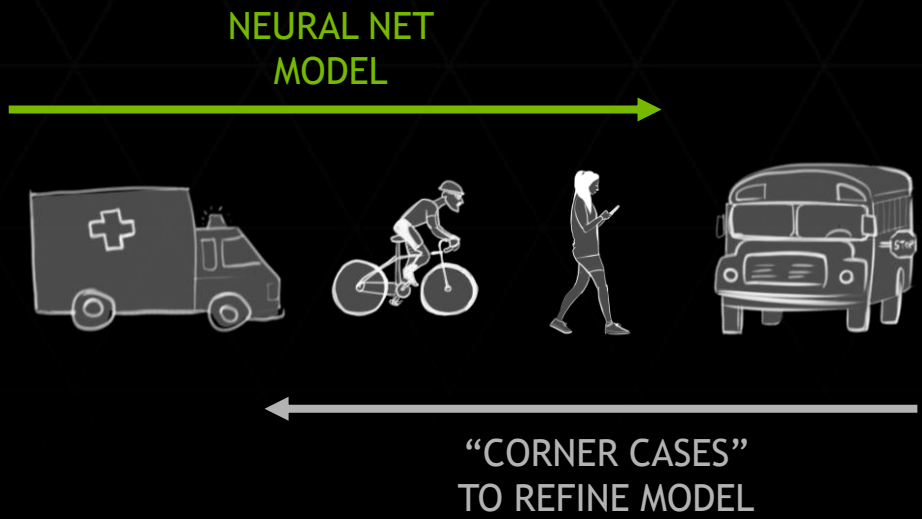
## OverFeat for Classification



# CARS THAT SEE BETTER



NVIDIA TESLA SUPERCOMPUTER  
FOR TRAINING

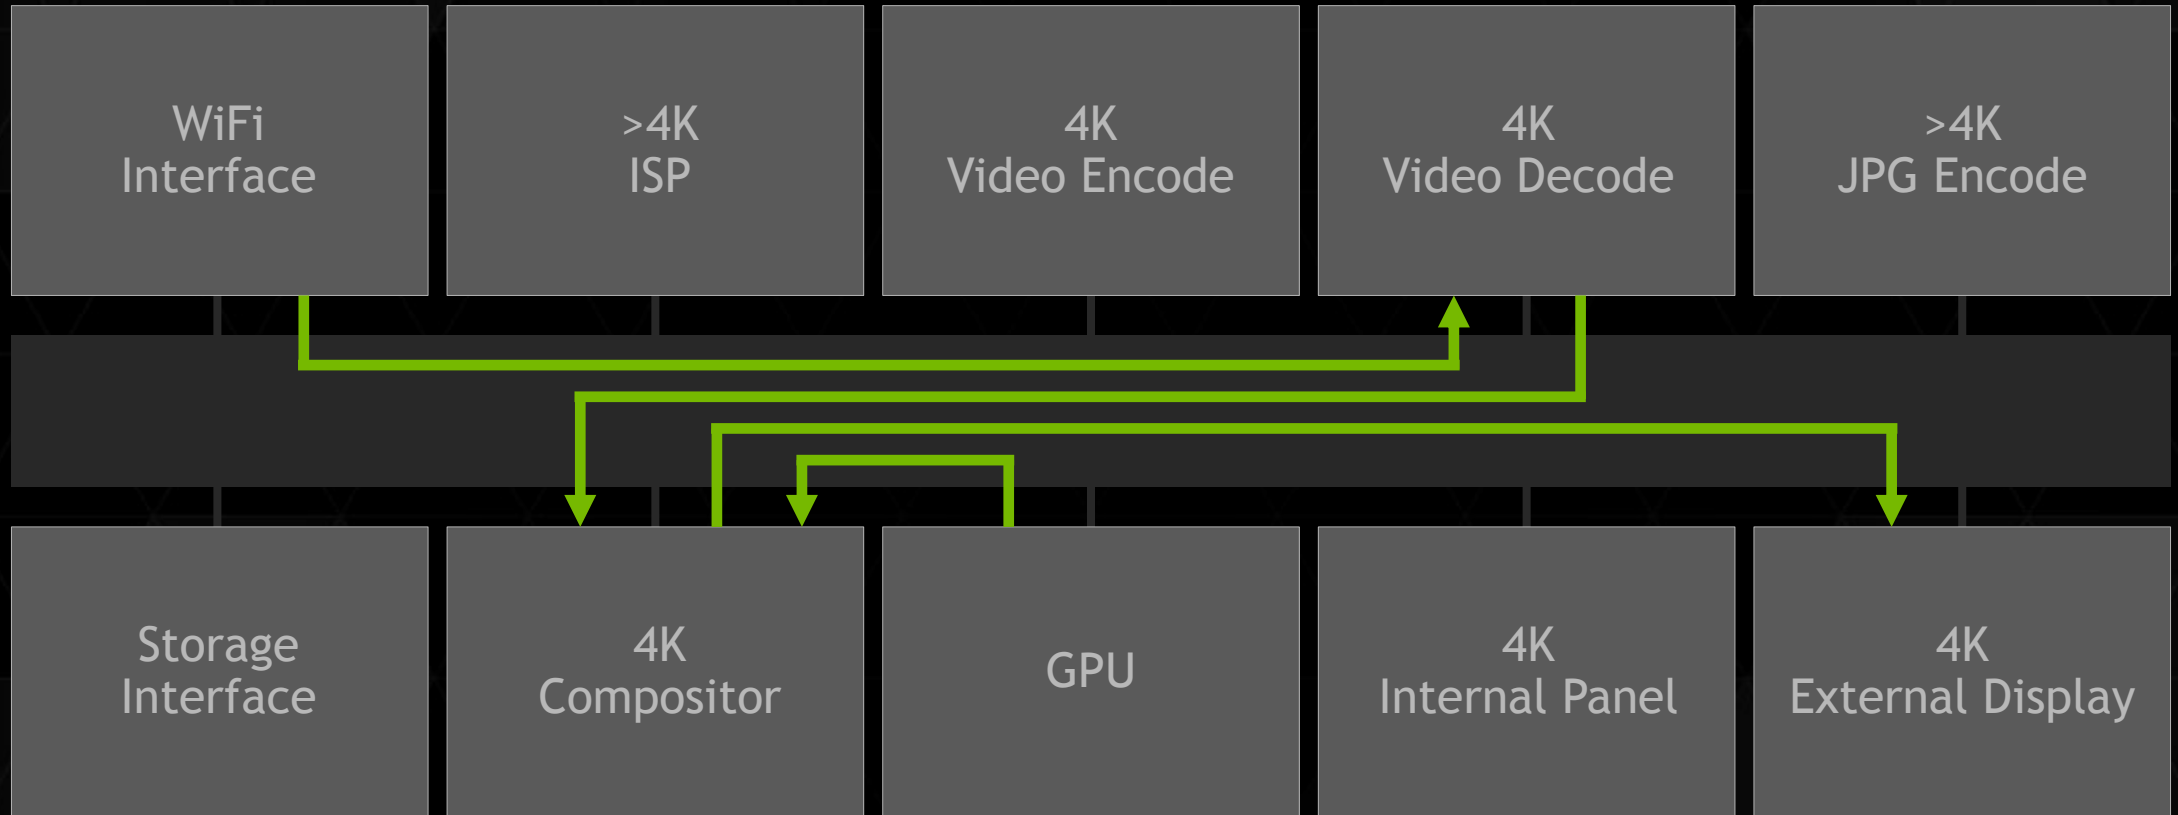


NVIDIA TEGRA SUPERCHIP  
FOR CLASSIFICATION



TEGRA X1  
4K VIDEO


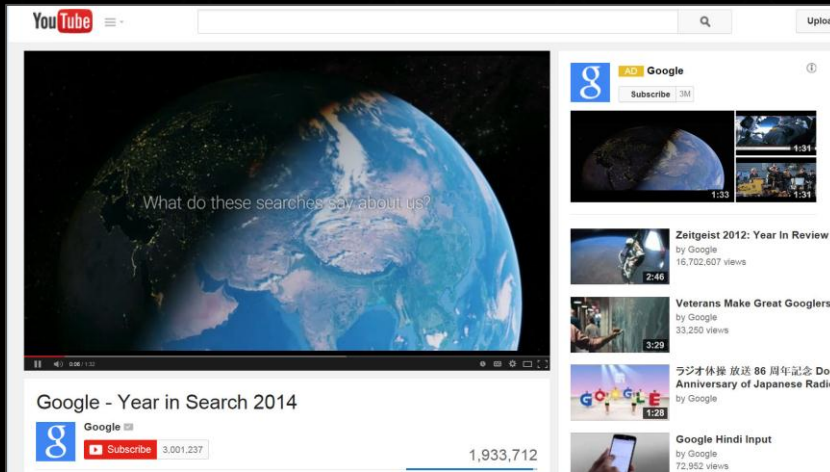
# END-TO-END 4K PIPELINE



60 fps 4K STREAMING

# THE FULL 4K EXPERIENCE

	TEGRA X1	OTHERS
60fps 4K VP9 & H.265	✓	✗
10-bit 4K H.265	✓	✗



# Tegra X1 – Video, Memory, Display Specs

DESCRIPTION	TEGRA K1	TEGRA X1	TEGRA X1 vs TEGRA K1
Video Decode			
VP9	-	2160p 60	New
VP8	2160p 30	2160p 60	2X
H.264	2160p 30	2160p 60	2X
H.265	-	2160p 60 (10-bit)	New
JPEG	120MPix/s	600MPix/s	5X
Video Encode			
VP8	✓	✓	
H.264	✓	✓	
H.265	-	2160p 30	New
JPEG	120MPix/s	600MPix/s	5X
Video Image Compositor	800MPix/s	4000MPix/s	5X
Internal Display	3200x2000 @60Hz	2160p @60Hz VESA DSC Compression	1.4X New
External Display	2160p 30 HDMI 1.4b HDCP 1.4	2160p 60 HDMI 2.0 HDCP 2.2	2X New New
Memory	LPDDR3 14.9GB/s (LP3)	LPDDR3, LPDDR4 25.6GB/s (LP4)	New 1.7X



**TEGRA X1  
SHIELD CONSOLE /  
ANDROID TV**

# BATTLE FOR THE LIVING ROOM

the 2. 202 on. SIDE 1

ed  
are  
t  
eir  
she  
that  
es to

**HBO NOW  
STAND-ALONE  
SERVICE  
SHAKES UP  
TV INDUSTRY**

A  
m  
age  
F  
Lyl  
in  
lin

Los Angeles Times

Apple is Out to Blow Up  
the Cable TV Model

**Bloomberg**

in  
the  
as a  
sation  
g was  
ket to  
e, and  
would  
e

The media industry is  
racing toward an  
Internet-TV future at  
a breathtaking pace

WALL STREET JOURNAL




# CHANNELS Become APPS



*“We are only at the beginning of this app driven evolution.”*

– Chris Louie  
VP @ Nielsen, Nov 2014

# ANDROID TV MOMENTUM



## Recently Announced

HBO NOW™

sling  
TELEVISION

HBO GO™

FXNOW

CBS

FOXNOW™

CBS  
SPORTS

FOX  
SPORTS

CBS  
NEWS

twitch

vudu  
HD Movies

Live  
Channels



# SHIELD

“The flagship of Android TV”



TEGRA X1: MOST ADVANCED MOBILE PROCESSOR

CONNECTED TO A WORLD OF CONTENT IN 4K

AMAZING GAMES

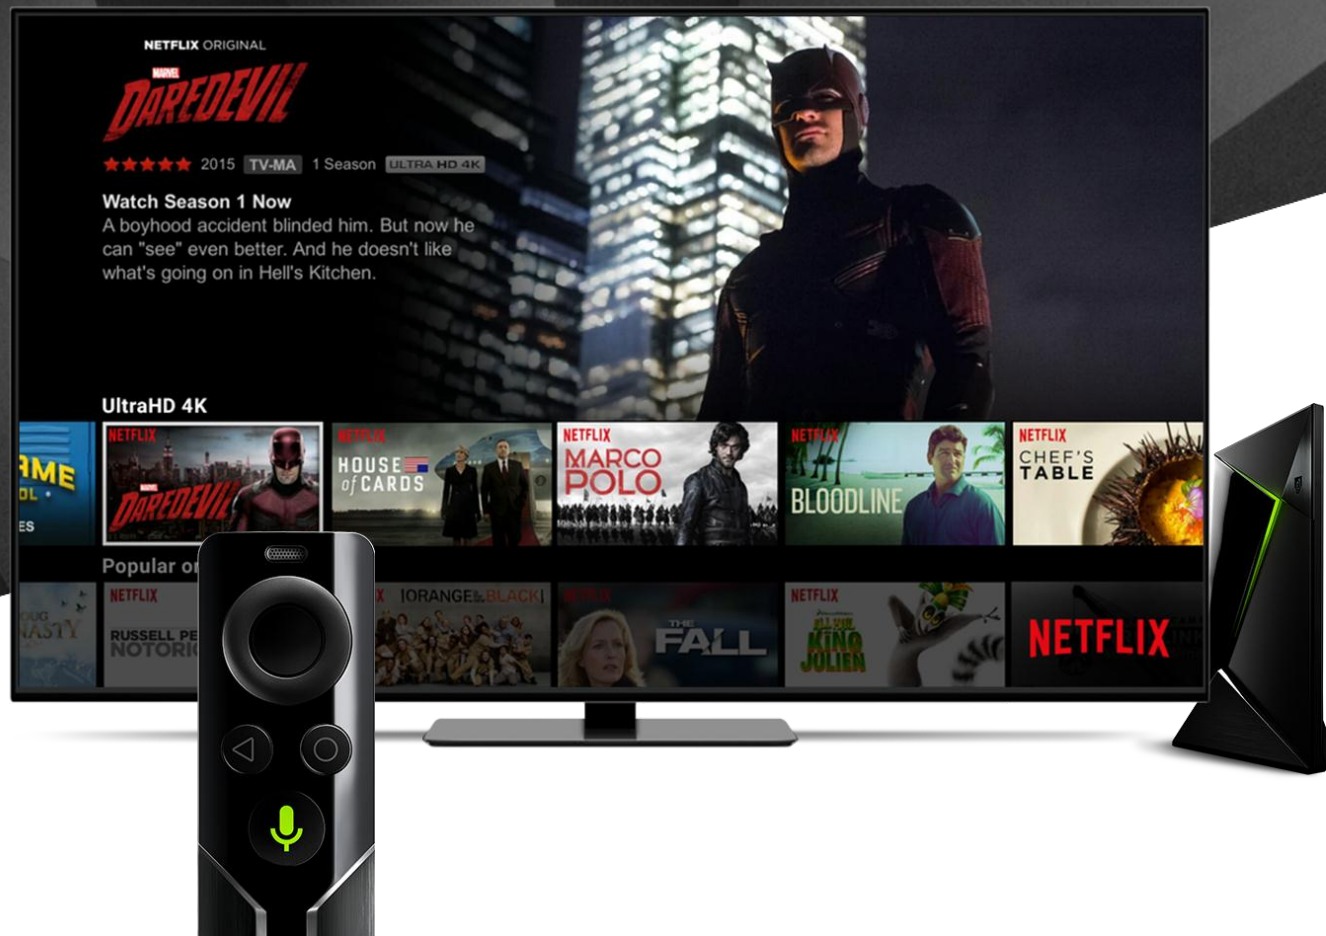
*“SHIELD is powered by Tegra X1, meaning it massively overpowers rivals like Fire TV and Apple TV.”*

*— Tech Radar*



*Data: GFXBench T-Rex off-screen*

# WITH 4K MOVIES AND SHOWS.



# WITH Amazing GAMES.



# WITH a world of apps.



# The FLAGSHIP OF ANDROID TV



## Best Performance

Tegra X1, 256 CORE NVIDIA GPU, 64-BIT CPU, 3GB

## Always Up-to-Date

NVIDIA OTAs

## Voice Search

## Private Audio



## Most Storage Options

MicroSD slot  
USB 3.0 x2

## Fastest Connectivity

Gigabit Ethernet  
802.11ac Wi-Fi



- 1 Processor and memory architecture
- 2 Maxwell GPU advances
- 3 Computer vision
- 4 4K video
- 5 Shield Android TV / console
- 6 **Demos**