



Whitepaper

NVIDIA Tegra K1

A New Era in Mobile Computing

Table of Contents

Introduction	3
NVIDIA Tegra K1 – A New Era in Mobile Computing	4
NVIDIA Kepler – The Most Advanced GPU comes to Mobile	7
Kepler Graphics Architecture in Tegra K1	10
Tessellation	11
Bindless Textures	12
Compute Shading	13
Kepler GPU Power Optimizations	15
LightSpeed™ Technology – High Performance Memory Architecture	16
Tegra K1 – A First Class Gaming Platform	17
Enabling Developers with World Class Development Tools	19
Tegra K1 – World’s First Mobile Processor to support Unreal Engine 4 Game Engine	20
Tegra K1 – The Leading Mobile GPU Compute Architecture	22
Next Generation Dual ISP Core	22
Conclusion	24
Document Revision History	25

Introduction

Smartphones, tablets, and clamshells are becoming primary devices of choice for a variety of activities such as reading email, taking pictures, playing games, interacting with social networks, and even creating new content. Automotive navigation, infotainment, and driver assist systems demand advanced visual computing capability and will be included in most automobiles in the near future. Responding to the growing performance needs of mobile use cases, newer mobile processors include advanced multi-core CPUs built on latest architectures and powerful GPU subsystems that enable new visual experiences.

Visual computing is the next frontier for mobile devices and automobiles. In addition to delivering rich graphical user interfaces, realistic 3D gaming, high resolution display outputs, and speedy Web page rendering, powerful GPUs are critical in enabling the next generation of mobile applications such as gesture and object recognition, motion tracking, computational photography, and augmented reality. These advanced visual computing applications require GPUs that not only deliver tremendous graphics performance and parallel computing performance, but also support the latest graphics and compute APIs (Application Processor Interfaces). Automotive navigation systems now include large, high-resolution screens with 3D Google Earth rendering, while driver assist systems perform object tracking and composite multiple video camera inputs, and entertainment systems deliver high-quality HD video to multiple screens.

The mobile gaming industry continues to grow rapidly and the visual quality of mobile games continues to improve with each new generation of mobile GPU subsystem. While casual games such as Angry Birds and Temple Run continue to be very popular, games such as Dead Trigger 2, Infinity Blade 2, Horn THD, and others are delivering great graphics and immersive gameplay. However, for mobile games to rival the visual quality and rich gaming experience of console and PC games, mobile processors require GPUs that are architected to deliver substantial graphics performance, support the latest API specifications such as OpenGL 4.4 and DirectX 11.2¹, and yet be highly power efficient to fit within mobile device power and thermal envelopes.

NVIDIA's latest and most advanced mobile processor **Tegra® K1**, creates a major discontinuity in the state of mobile graphics by bringing the powerful **NVIDIA Kepler™ GPU** architecture to mobile, and delivering tremendous visual computing capabilities and breakthrough power efficiency. Tegra K1 is offered in two pin-to-pin compatible versions. The first version uses a 32-bit quad-core, 4-PLUS-1 ARM Cortex A15 CPU. The second version uses a custom, NVIDIA-designed **64-bit dual Super Core CPU**. This CPU (codenamed "**Denver**") delivers very high single-thread and multi-thread performance. It is based on the **ARMv8** architecture, which brings the energy-efficient heritage of ARM processor technology to 64-bit computing. Both versions of Tegra K1 deliver stunning graphics and visual computing capabilities powered by the 192-core NVIDIA Kepler GPU.

¹ Kepler supports DX11.2 with Hardware Feature Level 11_0. For more details see http://nvidia.custhelp.com/app/answers/detail/a_id/3196/~/fermi-and-kepler-directx-api-support

With two decades of GPU industry leadership, the NVIDIA Tegra K1 mobile processor delivers the performance demanded by next-generation PC- and console-class mobile games, modern user interfaces, advanced visual computing applications, and high-resolution 4K displays, while delivering exceptional power efficiency required to bring the above capabilities to mobile devices.




Figure 1 Tegra K1 available in 32-bit and 64-bit versions that are pin compatible

NVIDIA Tegra K1 – A New Era in Mobile Computing

NVIDIA is at the forefront of the mobile computing revolution and has continually raised the bar on the performance and capabilities of mobile processors. NVIDIA launched Tegra 2, the world's first dual core mobile processor, Tegra 3 the world's first 4-PLUS-1 quad-core mobile processor, and Tegra 4 the world's first 4-PLUS-1 quad-core Cortex A15 processor while continuously raising the bar on graphics performance of mobile processors. NVIDIA Tegra mobile processors have enabled mobile devices to deliver snappy responsiveness, faster Web page load times, and visually rich gaming experiences while delivering long battery life through innovations such as the **4-PLUS-1™** CPU architecture, **PRISM** and **DirectTouch™** technologies.

The 32-bit and 64-bit versions of the NVIDIA Tegra K1 mobile processor are designed from the ground up to create a major discontinuity in the capabilities of mobile processors, and delivers the industry's fastest and most power efficient implementation of mobile CPUs, PC-class graphics and visual computing capabilities.

This whitepaper focuses primarily on the 32-bit version of Tegra K1. An upcoming paper will discuss in detail the 64-bit version of Tegra K1

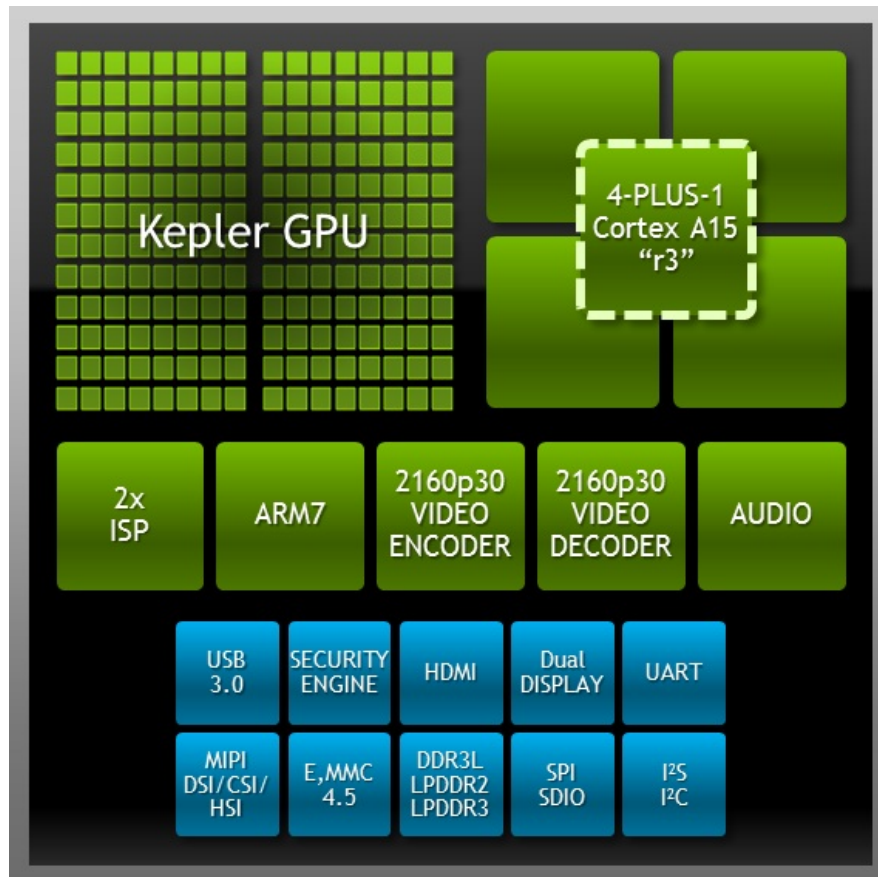


Figure 2 NVIDIA Tegra K1 Mobile Processor (32-bit version)

Some of the key features of the Tegra K1 SoC (System-on-a-Chip) architecture are:

- **4-PLUS-1 Cortex A15 "r3"** CPU architecture that delivers higher performance and is more power efficient than the previous generation.
- **Kepler GPU architecture** that utilizes 192 CUDA cores to deliver advanced graphics capabilities, GPU computing with NVIDIA CUDA 6 support, breakthrough power efficiency and performance for the next generation of gaming and visual computing applications.
- **Dual ISP Core** that delivers 1.2 Giga Pixels of raw processing power supporting camera sensors up to 100 Megapixels.
- **Advanced Display Engine** that is capable of simultaneously driving both the 4K local display and a 4K external monitors via HDMI
- Built on the TSMC **28 nm HPM** process to deliver excellent performance and power efficiency.

NVIDIA Tegra 4 was the first mobile processor to use a 4-PLUS-1 quad-core ARM Cortex A15 CPU architecture and **variable Symmetric Multiprocessing (vSMP)** technology to intelligently use the high performance quad-core A15 CPU complex for performance intensive tasks, and switch to the power

optimized “Battery Saver” A15 CPU core to handle low performance tasks and thus extend battery life. More details about NVIDIA’s vSMP technology and its benefits can be found [here](#)

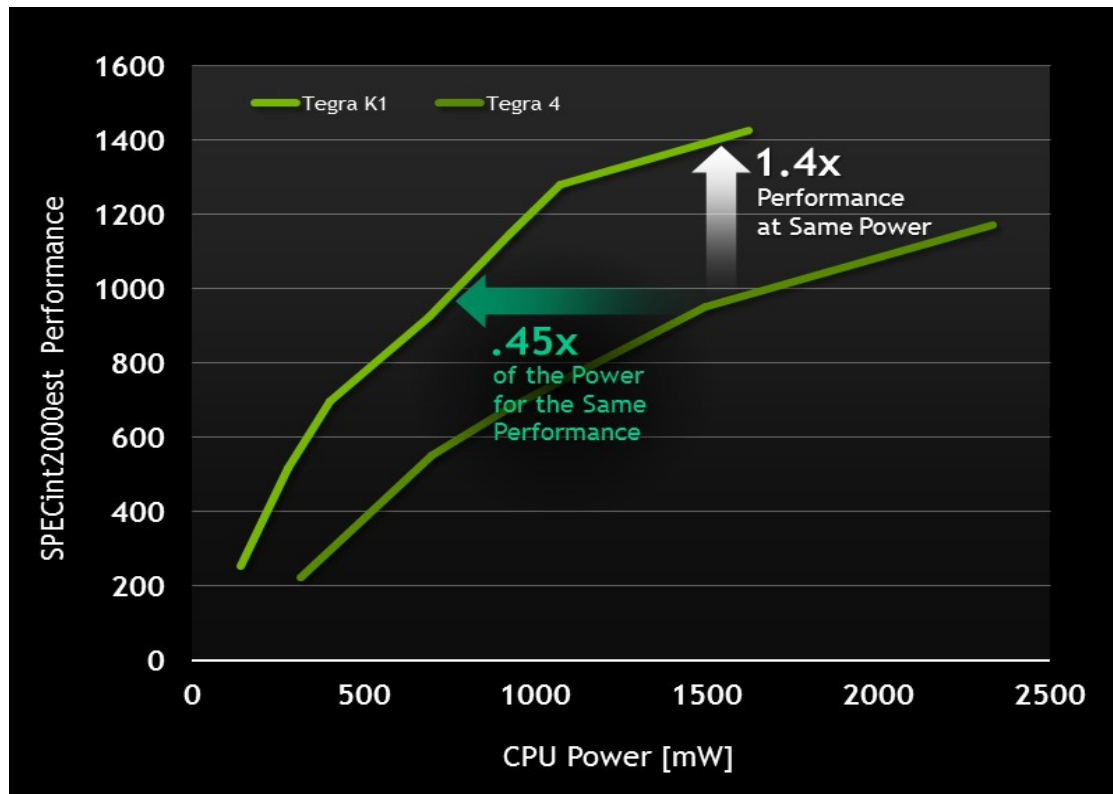


Figure 3 Tegra K1 Delivers higher CPU Performance and Power Efficiency²

The performance and power efficiency of the 4-PLUS-1 quad-core A15 CPU architecture in Tegra K1 is further optimized to use half the power for the same CPU performance as that of Tegra 4. Tegra K1 is more power efficient than Tegra 4, and also delivers almost 40% higher CPU performance for the same power consumption as Tegra 4.

Tegra K1 is the first mobile processor to use the latest revision (r3p3) of the ARM Cortex A15 CPU architecture that has several performance and power optimizations. In addition, ASIC-level optimizations based on learnings from Tegra 4 along with the use of the 28 nm HPM process helps Tegra K1 deliver exceptional CPU performance and power efficiency.

² Power and performance measured on Tegra K1 (32-bit) and Tegra 4 development systems while running SPECint2000 benchmark. SPECint2000 is a retired benchmark. Results are presented as “estimated” because SPECint2000 was run on pre-production Tegra K1 (32-bit) hardware. Final customer systems are expected to yield similar results. SPEC® and SPECint® are registered trademarks of the Standard Performance Evaluation Corporation. For more information about SPECint2000, see www.spec.org/cpu2000

NVIDIA Kepler – The Most Advanced GPU comes to Mobile

One of the most complex processors ever created, the GPU is the engine behind state-of-the-art computer graphics and energy efficient computing. NVIDIA's latest GPU architecture, codenamed Kepler, is the world's most advanced GPU architecture and powers some of the fastest and most power-efficient gaming systems, workstations, supercomputers, and cloud gaming servers.

Kepler-based GPUs such as the GeForce® GTX™ Titan and GTX 780M are the engines behind some of the world's highest performing gaming desktop and laptop PCs, respectively, delivering cutting-edge graphics performance for games such as Crysis 3, Assassin's Creed IV: Black Flag, Batman: Arkham Origins, and others. The Kepler-based NVIDIA Quadro® K6000 GPU powers many high-performance workstation systems that are used for 3D visualization and design, medical imaging, movie special effects, and numerous other scientific and engineering usages.

Kepler-based NVIDIA Tesla® GPUs used for High Performance Computing (HPC) power some of the world's fastest supercomputers such as the Titan supercomputer at Oakridge National Labs, and also the most power-efficient supercomputers such as the Eurora System at Cineca Supercomputing Center in Italy. Many scientific researchers, energy engineers, and financial analysts throughout the world also use Kepler-based Tesla GPUs in their own personal and workgroup HPC systems.

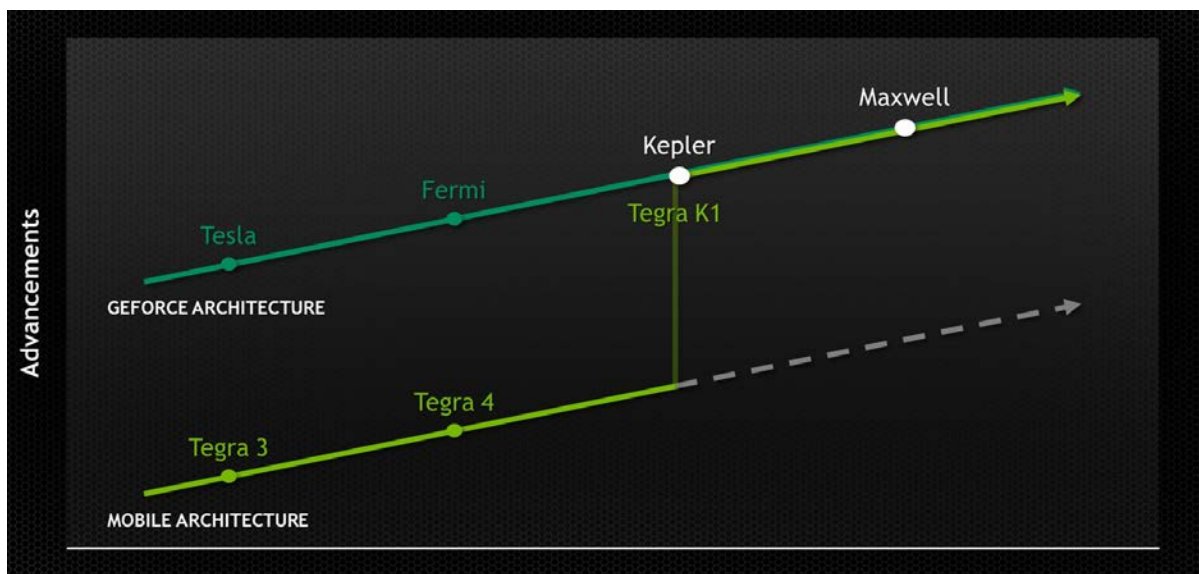


Figure 4 Mobile Kepler in Tegra K1 brings desktop class graphics architecture to Mobile

The Kepler GPU in Tegra K1 is built on the same high performance, energy efficient Kepler GPU architecture that is found in these higher-end systems, and is an industry breakthrough for mobile graphics and computing. Kepler delivers the most advanced graphics for mobile gaming, and is the first modern mobile GPU capable of supporting all the GPU compute APIs.

The Kepler architecture boasts several key features such as the advanced **Streaming Multiprocessor (SMX)** for higher performance and power efficiency, **Polymorph Engine v2.0** for improved **Tessellation** performance, **bindless textures** enabling up to one million simultaneous textures and higher efficiency, **FXAA™** and **TXAA™** anti-aliasing techniques for smoother visuals, and higher performance, **PhysX™** for realistic collision, explosion, water, and smoke effects, and a new display engine for the next generation of 4K displays. A detailed explanation of these features and their benefits can be found in the Kepler architecture whitepapers titled "[NVIDIA GeForce GTX 680](#)" and "[NVIDIA's Next Generation CUDA Compute Architecture](#)".

In addition Kepler will be the first mobile GPU that supports key API specifications such as **OpenGL 4.4**, **DirectX 11.2**, and **CUDA™ 6³**.

	TEGRA K1 Kepler Graphics	GEFORCE TITAN
OpenGL ES 3.0	✓	✓
OpenGL 4.4	✓	✓
DX11	✓	✓
Tessellation	✓	✓
CUDA 6.0	✓	✓

Figure 5 Tegra K1's Kepler GPU supports the same features as the most powerful desktop gaming GPUs

The Kepler GPU delivers the graphics features, rich APIs, and compute architecture of its desktop counterpart, and has additional power optimizations for mobile usage.

The Kepler GPU in Tegra K1 is a significant milestone in the history of computing and computer graphics and will drive a revolutionary change in mobile visual computing.

³ Tegra K1 is capable of OpenCL 1.2. It will be supported based on customer needs.



Figure 6 Kepler is a major discontinuity in Mobile Graphics



Figure 7 Tegra K1 brings never before seen visual realism to mobile devices

Kepler Graphics Architecture in Tegra K1

The architecture of the Kepler GPU in Tegra K1 is virtually identical to the Kepler GPU architecture used in high-end systems, but also includes a number of optimizations for mobile system usage to conserve power and deliver industry-leading mobile GPU performance. While the highest-end Kepler GPUs in desktop, workstation, and supercomputers include up to 2880 single-precision floating point CUDA cores and consume a few hundred watts of power, the Kepler GPU in Tegra K1 consists of 192 CUDA cores and consumes less than two watts⁴. Note that the Tegra K1 Kepler GPU has more cores than many entry-level to mainstream desktop GPUs of just a few years ago.

The Kepler architecture is organized in Graphics Processing Clusters (GPC), Streaming Multiprocessors (SMX), and memory controllers. For example, the Kepler based desktop GeForce GTX 680 GPU consists of four GPCs, eight SMX, and four memory controllers. The Kepler GPU in Tegra K1 consists of one GPC, one SMX unit and a memory interface. The Tegra K1 GPU includes four ROPs and has a 128KB L2 cache between the ROPs and the memory interface.

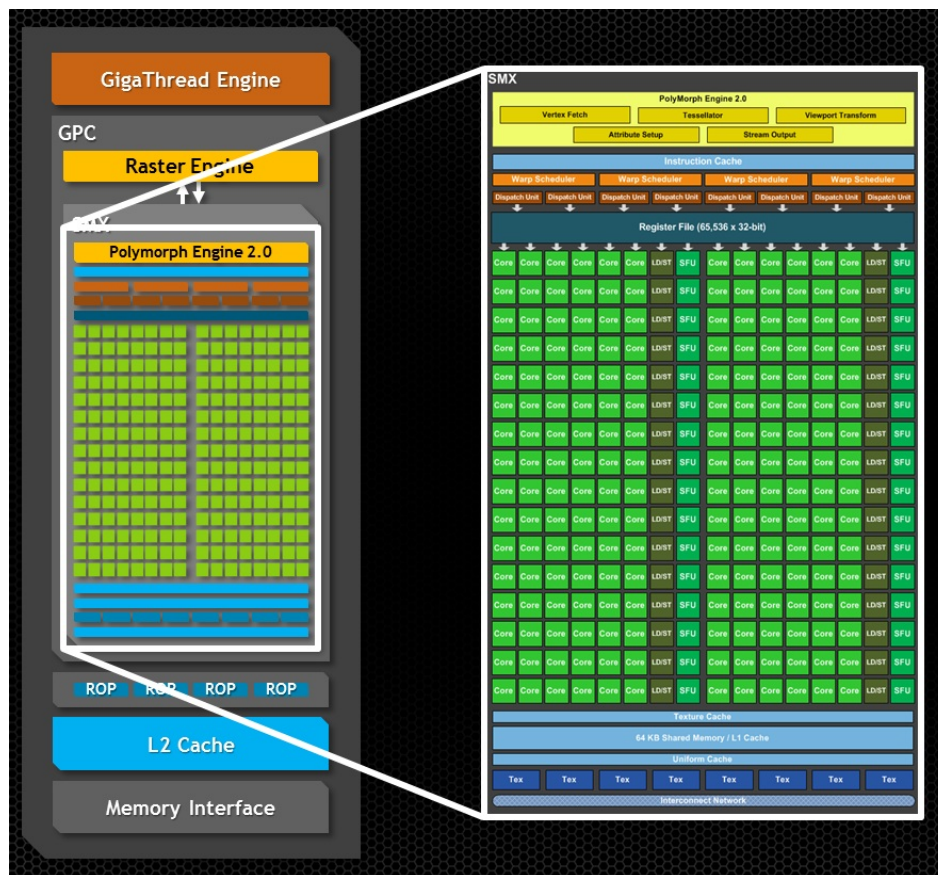


Figure 8 Full Kepler GPU (left) and Kepler SMX Unit (right) in Tegra K1

⁴ Average power measured on GPU power rail while playing a collection of popular mobile games.

The Kepler GPU in Tegra K1 supports the DX11.2 API with Hardware Feature Level 11_0, similar to the desktop Kepler GPUs, the latest OpenGL 4.4 specification, and NVIDIA CUDA 6, delivering several key features such as Tessellation, Bindless Textures, Global Illumination, deferred rendering, and advanced post-processing that will revolutionize graphics in mobile gaming.

Tessellation

Tessellation is one of the key features of OpenGL 4.x and DirectX 11.x that has profoundly influenced 3D graphics for PC gaming, and has increased the level of visual realism in PC games to being almost film-like. Mobile Kepler GPU is the first mobile GPU to support hardware tessellation, delivering breakthrough geometric complexity and visual realism to mobile games and applications. More details on how tessellation works can be found [here](#). Tessellation delivers more detailed terrains, character models, and environments.

OpenGL4.x hardware-based tessellation allows the GPU to dynamically generate additional polygons based on the user's viewpoint in a given scene. In contrast, OpenGL ES 2.0 requires pre-rendering up to fifty times more polygons to achieve the same visual quality and detail as that of hardware-based tessellation. Rendering these additional polygons in OpenGL ES2.0 is both performance intensive and power inefficient. Kepler's support of OpenGL4.x based hardware tessellation allows it to deliver more detailed scenes at much higher frame rates and lower power consumption.




Figure 9 Tessellation adds detail, depth and realism



Figure 10 Realistic Water and Terrain due to tessellation

Bindless Textures

In traditional GPU architectures, for the GPU to reference a texture, the texture had to be assigned a “slot” in a fixed-size binding table. The number of slots in that table ultimately limits how many unique textures a shader can read from at run time.

With bindless textures in Kepler, the shader can reference textures directly in memory, making binding tables obsolete. This effectively eliminates any limits on the number of unique textures that can be used to render a scene. As a result, many more different texture materials can be used to increase the texture detail in a game. Another benefit of bindless textures is the reduced driver and application overhead and lower CPU utilization.

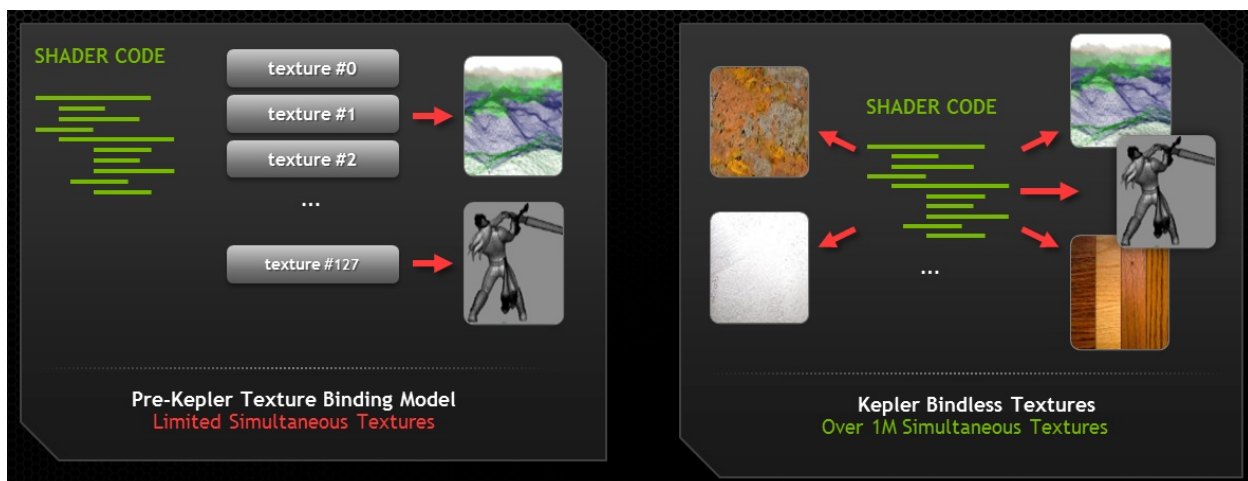


Figure 11 Bindless textures delivers higher texture detail and lower CPU utilization in games

Compute Shading

Along with tessellation, Compute Shading is a key feature of OpenGL 4.3 and 4.4, DirectCompute, and CUDA that helps deliver advanced graphics effects to games and 3D applications. Unlike other shader stages such as vertex and pixel shading that have well-defined set of inputs, Compute Shading is used for computing arbitrary information. Compute Shading is extremely bandwidth efficient and power efficient. Similar to other parallel programming systems such as CUDA, DirectCompute, and OpenCL, OpenGL 4.3 and 4.4 Compute Shaders naturally express parallelism and are very GPU-centric.

Compute shaders are primarily used to deliver a variety of advanced graphics effects such as particle physics effects, realistic fluid behavior, crowd simulation, global illumination, and many post-processing effects such as edge-sharpening, various types of ambient occlusion, blur, Depth-of-Field, and others.

Compute Shading is used for Global illumination, an important graphical effect that delivers more realistic lighting effects. Global illumination not only factors in direct light in a scene, but also light rays that are reflected from other surfaces in the scene. OpenGL Compute Shading does the complex calculations to account for light rays bouncing off surfaces that have varying reflective, refractive and absorption characteristics. Global illumination and ambient occlusion can add great depth and richness to a scene.



Figure 12 Compute Shading-based Global Illumination effects delivers rich depth to scenes

When using Compute Shading for advanced post-processing, which is applied after the primary “3D” rendering of a scene is complete, more sophisticated effects can be generated. For example, using tri-diagonal solvers similar to those used in film production, more realistic depth of field effects can be produced that do not display a general blurriness or color bleeding associated with less sophisticated techniques.

The NVIDIA Faceworks graphics demo was created to illustrate the visual realism and richness that is possible with tessellation, global illumination, post-processing, sub-surface scattering, and other advanced graphics effects.

The NVIDIA Faceworks demo gives a glimpse of the realism that will now be possible in mobile games. Originally created to demonstrate the capabilities of the NVIDIA GeForce TITAN GPU, porting the Faceworks demo to mobile was possible only because the Kepler GPU in Tegra K1 shares the same architectural features as the more powerful desktop Kepler GPUs. While some features were scaled back versus the desktop version, the Kepler GPU in Tegra K1 still leveraged OpenGL 4.x features such as tessellation, geometry shaders, compute shaders, TXAA, and other post processing effects to deliver almost life-like rendering of the human face.



Graphics on current OpenGL 3.0 ES based mobile SoCs



Revolutionary graphics on Tegra K1 using OpenGL 4.x

Figure 13 Revolutionary graphics in Tegra K1 delivers unrivaled realism in graphics

Kepler GPU Power Optimizations

The Tegra K1 Kepler GPU architecture is not only optimized for high performance, but also for efficient power utilization. It includes the same power-efficient design found in its desktop counterpart and adds several new optimizations that help deliver very low power consumption in both active use and idle conditions. Several key features help enable higher power efficiency such as hierarchical on-chip Z cull, primitive culling, Early Z culling, Texture, Z, and color compression, and a large unified L2 cache that significantly decreases accesses to power hungry off-chip memory.

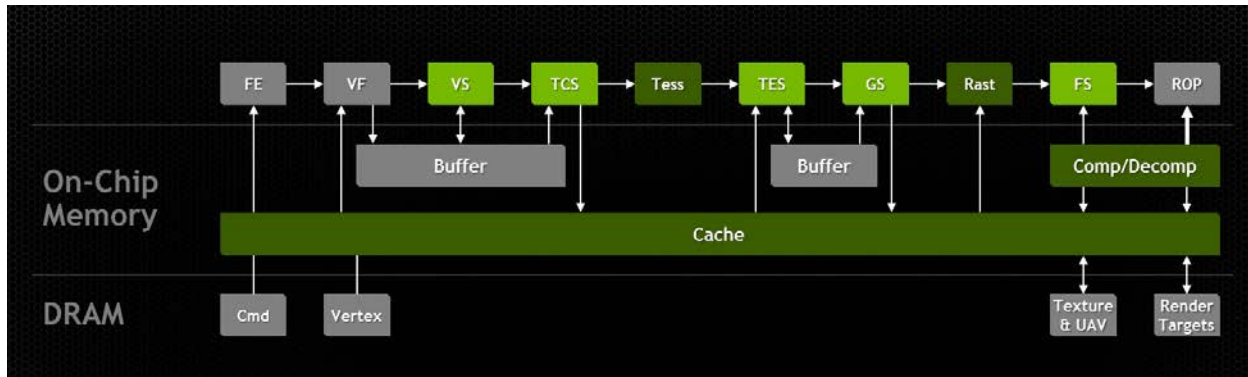


Figure 14 On-chip unified L2 cache (the big green Cache block) helps reduce off-chip memory accesses and improves power efficiency




Figure 15 Tegra K1 consumes significantly lower power when configured to deliver the same peak performance of competing SoCs⁵ (Performance is Y-axis and Power is X-axis)

⁵ Based on measured data on GFXbench 3.0 GL Gold, Manhattan test at 1080p off-screen resolution

LightSpeed™ Technology – High Performance Memory Architecture

The Kepler GPU in Tegra K1 supports all three texture compression formats: DXT, ETC and ASTC. Compression plays an important role in reducing both off-chip memory traffic and power consumption. Even for non-gaming applications such as the Android User Interface and Web browsers, frame buffer compression technologies used in the Tegra K1 Kepler GPU reduce power consumption by 40-70% in many scenarios measured by NVIDIA engineering.

NVIDIA LightSpeed frame buffer compression technology has been developed and refined over the years on NVIDIA desktop and mobile GPUs, and is very effective in reducing memory bandwidth and decreasing power utilization by saving power-hungry, off-chip memory accesses. Further optimized and improved for Tegra K1 compared to Tegra 4, NVIDIA LightSpeed technology is extremely effective for mobile user interfaces. For example, LightSpeed's lossless compression is used on the background layer of the Android UI, the icon layer of the Android application menu and even on content displayed within a Web browser window. LightSpeed technology reduces off-chip system memory accesses and delivers almost 40-70% percent lower memory bandwidth and consequently delivers considerable power savings.




Figure 16 – LightSpeed's lossless compression delivers higher available bandwidth and saves power

Along with the above power saving architectural features, the Tegra K1 Kepler GPU includes several low level optimizations to reduce both idle and dynamic power consumption. Features such as Rail Gating, Power Gating, and Clock gating identify blocks of the GPU core that are idle, and turn off both clock and voltages sources to reduce idle power consumption of these blocks. By defining activity regimes within the GPU core, power management algorithms can dynamically gate clocks and or reduce clock frequencies to reduce dynamic power consumption. Optimized routing of data paths and interconnects further help reduce power consumption.

Tegra K1 – A First Class Gaming Platform

Mobile gaming is in a rapid growth phase and at its current pace, by 2015, mobile game revenues are set to overtake both PC and console game revenues. Almost 74% of revenue in the Google Play Store is derived from gaming apps⁶, and studies show that almost 70% of consumer tablet usage time is spent playing games⁷. Developers are responding to the skyrocketing growth in mobile gaming and more developers are spending time developing games for mobile devices than for PCs and consoles.

Recognizing the rapid growth in mobile gaming revenues, developers are beginning to launch games with advanced graphics effects, richer and more engaging storylines, and immersive gameplay experience on mobile devices. Developing such games require significant amount of technical manpower, game development tools, complex game engines, and graphics processors that are not only powerful, but also support advanced graphics APIs and feature sets.




Figure 17 Mobile Gaming Revenues are growing at a rapid pace⁸

PC and console gaming platforms are supported by well-defined development tools, powerful game engines, and established APIs such as DX11.x and OpenGL4.x that are well understood by game developers enabling them to develop games on either platform, or easily port to the other platform.

On mobile platforms, developers have found it difficult to develop advanced games, or port existing console/PC games to mobile architectures for several reasons. First, mobile gaming platforms deliver limited graphics performance compared to consoles and PCs. Second, mobile SoCs currently support older, mobile-specific graphics APIs called OpenGL ES2.0 and ES 3.0 that do not support advanced

⁶ Reported by App Annie on U.S. Revenues on Google Play in December 2012

⁷ Flurry Analytics Report

⁸ Data as reported by DFC

features such as hardware tessellation, Geometry Shaders, and Compute Shaders that are found in DX11.x and OpenGL4.4. Third, a game development ecosystem that exists for PCs and consoles, including powerful game engines, advanced debuggers and profilers, and a rich set of libraries for visual effects and physics has been lacking.

Tegra K1 is the first mobile processor to support the very popular Unreal Engine 4 (UE4) game engine, as we described in more detail below, allowing easy porting of state-of-the-art PC-class games to mobile.

Tegra K1 with its Kepler GPU core removes the performance and graphics feature-set barriers by delivering higher peak shader GFLOPs and total CPU horsepower than Xbox 360 and PS3 consoles, while supporting more advanced OpenGL4.4 and DX11.2 APIs found in enthusiast gaming GPUs such as the GeForce GTX 780. By supporting the same graphics APIs and rendering pipeline as more powerful desktop GPUs, developers will now be able to develop a game on a console or PC platform and use the same codebase to port the game to a Tegra K1-powered mobile platform. While games developed for higher-end desktop GPUs may require some performance profiling and feature tuning before they can be run on Tegra K1 powered mobile devices, a vast majority of Xbox 360 and PS3 titles could be easily ported over to Tegra K1 devices with minimal effort, delivering similar visual quality and gameplay experience.

	Xbox360	PS3	Tegra K1
GPU Features	DX9	DX9	DX11
GPU Horsepower	240	192	365
CPU Horsepower	3600	1200	5612
Power	100W	100W	5W

Figure 18 Tegra K1 delivers much higher graphics processing power than current generation consoles⁹

⁹ GPU Horsepower is based on Peak Fragment Shader GFLOPS of each platform. PS3 GPU-Cell 154 GFLOPS FP32 not included. CPU Horsepower is the estimated SPECint2000 performance multiplied by the number of CPUs. Data is based on internal measurements. Peak GFLOPS and SPECint2000 performance of consoles are estimates and may have +/-20% margin of error. SPU's not included in the calculation of CPU Horsepower for PS3



Figure 19 PC and Console Games on Tegra K1 powered devices.

Enabling Developers with World Class Development Tools

The Tegra K1 Kepler GPU is based on the same desktop Kepler GPU architecture used in powerful desktop GPUs and is capable of running the same set of rich developer tools, game libraries, debugging software, and profiling tools that have been developed and optimized over the years for desktop GPUs.

NVIDIA GameWorks research group consists of over 300 world class graphics scientists and engineers working at the intersection of art and science to create powerful new visual effects libraries, development tools, new graphics algorithms, game engine optimizations, and SDKs. Technologies developed by NVIDIA GameWorks group will continue to improve the visual realism of games, and empower game developers with the latest toolkits for world class game development. The GameWorks tools provide a seamless developer experience across GeForce, Quadro, Tesla and Tegra platforms.

NVIDIA's gaming heritage and over 20 years of experience in graphics and gaming is unrivaled in the mobile space. NVIDIA has long standing relationships with all the major game engine developers, and is a key technical partner for all major game engines. NVIDIA has access to source code of popular game engines, and contributes regularly to the code base of these engines to make them run more efficiently on NVIDIA GPUs. Popular game engines such as Epic's Unreal Engine 4, Unity, Frostbite, CryEngine 3, id Tech, and Source are optimized to run on NVIDIA Kepler GPU architecture, and will run more efficiently on NVIDIA Tegra K1.



Figure 20 NVIDIA Development tools scale from GeForce to Tegra K1

Tegra K1 – World’s First Mobile Processor to support Unreal Engine 4 Game Engine

Unreal Engine 4 (UE4) is the latest version of Epic’s enormously popular game engine and a wide variety of PC and console games based on UE4 are expected to be launched in 2014. UE4 enables games with advanced lighting, a full post-processing pipeline for Bloom, Camera Effects, Depth of Field, PhysX, and host of other OpenGL4.4 and DX11 features. The Kepler GPU in Tegra K1 supports all the key features required to run Unreal Engine 4 such as compute shaders, geometry shaders, multiple render targets, cube map arrays and others. NVIDIA’s close technical partnership with Epic, coupled with Kepler’s support for the latest desktop graphics feature set, has enabled Epic to effortlessly get UE4 up and running on Tegra K1-powered development platform.

“With the introduction of this revolutionary processor, we can take absolutely anything that runs on PCs or high-end consoles and run it on Tegra. From here onward, I think we’re going to see the performance and feature gap between mobile and PC high-end gaming continue to narrow to the point where the difference between the platforms really blurs” -Tim Sweeney, Founder of Epic Games and Unreal Engine




Figure 21 UE4 based games deliver advanced visual effects on Tegra K1



Figure 22 UE4 based games on Tegra K1 takes mobile graphics to the next level

Developers who are developing UE 4 based games for next-generation game consoles and PCs can now also easily port these games to Tegra K1-powered mobile platforms.

NVIDIA Tegra K1 has created a major discontinuity in the capabilities of mobile graphics by delivering higher performance than current generation consoles, and enabling PC class features such as tessellation, advanced lighting, global illumination, post-processing and PhysX. Tegra K1 along with NVIDIA's powerful and rich set of development tools and developer support has effectively removed all barriers to bring visually compelling and rich PC and console-class gaming experiences to mobile.

Tegra K1 – The Leading Mobile GPU Compute Architecture

General Purpose GPU Computing (GPGPU) refers to the use of a GPU to greatly accelerate compute-intensive applications. GPGPU was pioneered by NVIDIA years ago, permitting the GPU to be directly programmed to run compute-intensive applications. Later, with the introduction of GeForce 8800 (G80) in 2006, NVIDIA's CUDA parallel programming platform standardized and refined GPGPU computing. CUDA has become the world's leading GPU computing platform used by millions of users for high-performance computing across a range of industries and sciences, including usage in many of the top supercomputers in the world. GPU computing delivers unprecedented levels of performance speedups by parallelizing application workloads and running them on the GPU.

NVIDIA Kepler GPU architecture is designed to deliver tremendous amounts of compute performance while being almost three times more power efficient than the previous generation NVIDIA Fermi GPU architecture. The Kepler Compute architecture together with NVIDIA's CUDA parallel programming platform delivers tremendous performance speedup not only for numerous high-performance computing applications, but also for applications such as speech recognition, live video processing, computer vision, augmented reality and of course computer gaming.

The NVIDIA Tegra K1 Kepler GPU is the world's first mobile GPGPU architecture that delivers considerable amounts of compute performance, and can support many different types of Compute Shaders used in games, enable real-time beautification of videos, or computation photography effects including real-time local tone mapping, and real-time video editing. Tegra K1 gives developers the flexibility to create compelling compute-based mobile applications.

Next Generation Dual ISP Core

The NVIDIA Tegra K1 processor includes a brand new, high performance dual-ISP core that supports up to a 100 Megapixel sensor, handling 4096 simultaneous focus points, and providing a throughput of 1.2 Gigapixels per second.

Each ISP core delivers significantly higher image quality through advanced features such as Spatial Variable Noise Reduction, Multi-bad pixel correction, high quality downscaling, flexible color space converters, and a state-of-the-art Area Processor. These features help deliver crisper images that have much lower noise, improving the performance of compute-based applications such as computer vision.

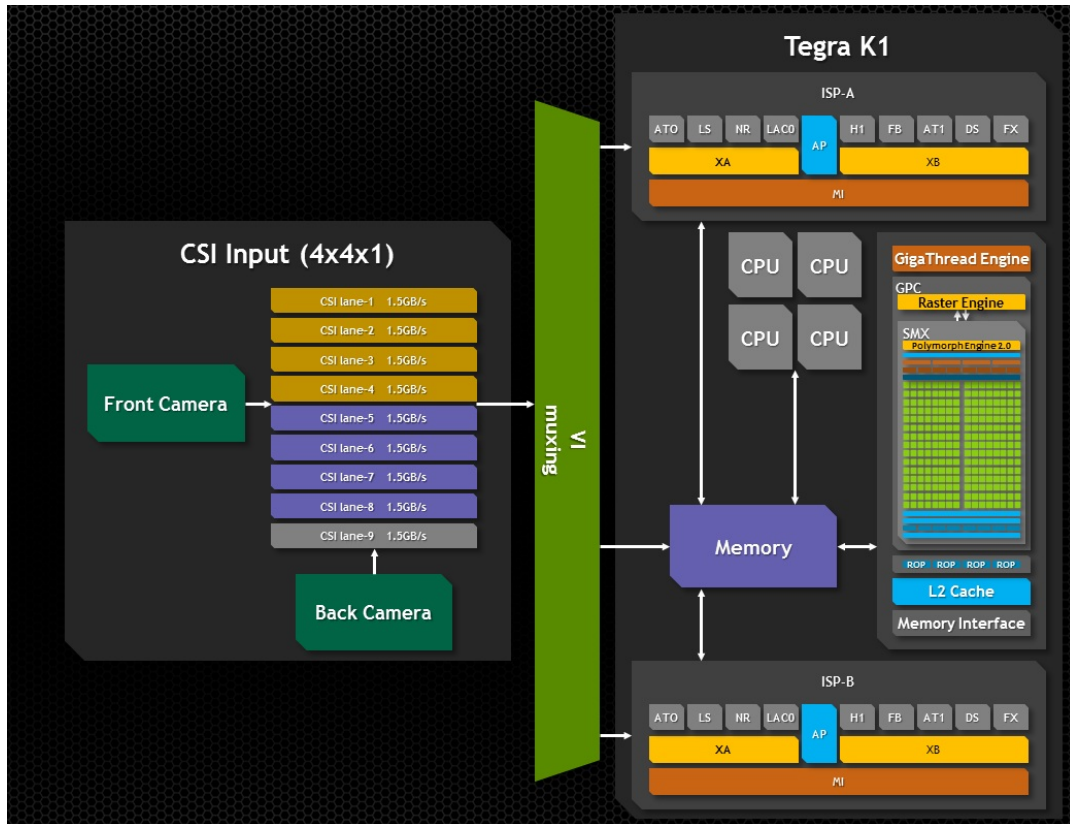


Figure 23 Powerful dual-ISP core architected to leverage the compute power of Kepler

The powerful ISP core uses up to 4096 focus points to deliver much higher performance for specialized focus functions such as auto-focus, focus-time, and low-light focusing .

The ISP core is architected for full GPGPU interoperability and seamlessly works with the Tegra K1 Kepler GPU to deliver advanced features.

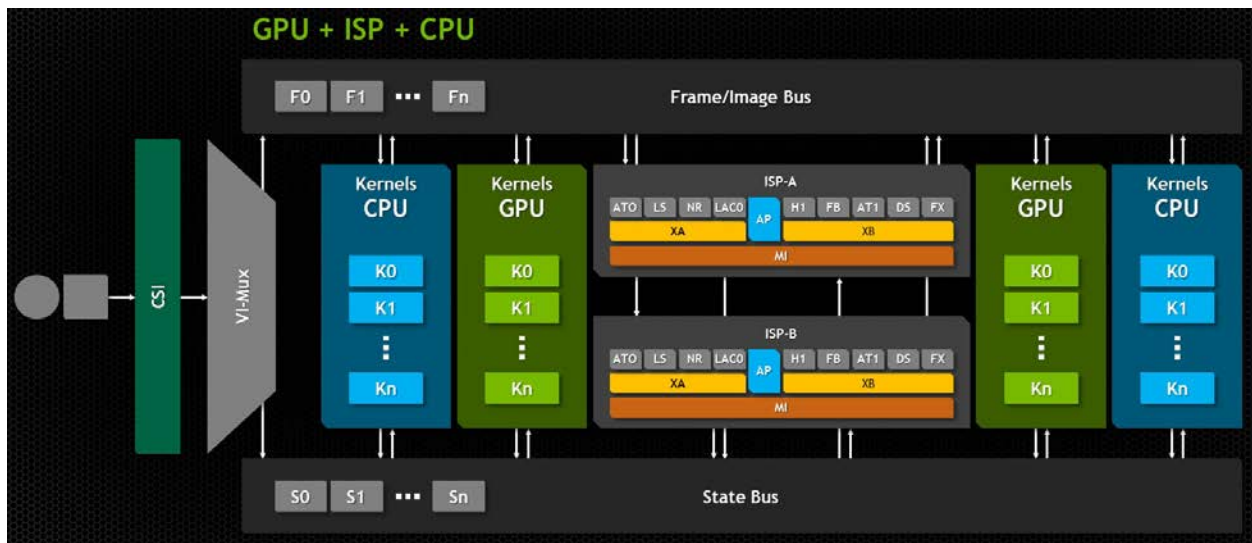


Figure 24 Camera Software Architecture

Conclusion

Smartphones, tablets, and automotive navigation, infotainment, and instrument clusters are quickly becoming our most frequently used visual computing systems. High-performance, power-efficient processors are required to deliver the types of mobile experiences users demand. The NVIDIA Tegra K1 processor creates a major discontinuity in the state of mobile visual computing by incorporating the industry's fastest and most power efficient GPU architecture.

Tegra K1 is NVIDIA's first SoC processor to include our high-performance Kepler GPU architecture. The 64-bit version of Tegra K1 brings very high single-threaded and multi-threaded CPU performance to mobile computing. PC class applications that required high performance and large amounts of memory can now be developed for Tegra K1 powered mobile devices. The 32-bit and 64-bit versions of Tegra K1 will power a wide breadth of devices including smartphones, tablets clamshells, All-in-One PCs, micro-consoles, mobile gaming devices, and automotive instrument panels.

The Kepler GPU delivers industry-leading mobile visual computing performance, exceptional power efficiency, and a rich set of desktop PC-class graphics and compute APIs. Because the Tegra K1 Kepler GPU architecture is virtually identical to the desktop Kepler GPUs, developers can easily develop or port games to run on PCs, consoles, and mobile devices without having to invest significant additional resources for each platform. The Tegra K1 Kepler GPU not only enables groundbreaking mobile graphics effects, it's the first mobile GPU to support the latest Unreal Engine 4 (UE4) game engine. NVIDIA's investment in GPU technology, game engines, and development tools can now be leveraged by Tegra K1 and help developers deliver compelling visual experiences on mobile devices.

NVIDIA Tegra K1 will usher in the next generation of new, exciting, high-quality mobile visual computing experiences powered by its Kepler GPU, NVIDIA's world-class development tools and utilities, and a graphical API feature-set unrivalled in the mobile industry today.

Document Revision History

- Release 1.0 – Initial Release
- Release 1.1 – Includes corrections of SPECint2000 naming and explanations of SPECint2000 benchmark results.

Notice

ALL INFORMATION PROVIDED IN THIS WHITE PAPER, INCLUDING COMMENTARY, OPINION, NVIDIA DESIGN SPECIFICATIONS, REFERENCE BOARDS, FILES, DRAWINGS, DIAGNOSTICS, LISTS, AND OTHER DOCUMENTS (TOGETHER AND SEPARATELY, "MATERIALS") ARE BEING PROVIDED "AS IS." NVIDIA MAKES NO WARRANTIES, EXPRESSED, IMPLIED, STATUTORY, OR OTHERWISE WITH RESPECT TO MATERIALS, AND EXPRESSLY DISCLAIMS ALL IMPLIED WARRANTIES OF NONINFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE.

Information furnished is believed to be accurate and reliable. However, NVIDIA Corporation assumes no responsibility for the consequences of use of such information or for any infringement of patents or other rights of third parties that may result from its use. No license is granted by implication or otherwise under any patent or patent rights of NVIDIA Corporation. Specifications mentioned in this publication are subject to change without notice. This publication supersedes and replaces all information previously supplied. NVIDIA Corporation products are not authorized for use as critical components in life support devices or systems without express written approval of NVIDIA Corporation.

Trademarks

NVIDIA, the NVIDIA logo, Chimera, Tegra, TegraZone are trademarks or registered trademarks of NVIDIA Corporation in the United States and other countries. Other company and product names may be trademarks of the respective companies with which they are associated.

Copyright

© 2013 NVIDIA Corporation. All rights reserved.

