



Whitepaper

# The Benefits of Multiple CPU Cores in Mobile Devices

## Table of Contents

Executive Summary.....	3
The Need for Multiprocessing.....	3
Symmetrical Multiprocessing (SMP) .....	4
Introducing NVIDIA Tegra 2: The world’s first mobile processor with a Dual-Core CPU .....	7
Dual-core ARM Cortex A9 Architecture .....	7
Benefits of Symmetrical Multiprocessing in Tegra 2 .....	9
Faster Web page load times.....	9
Lower Power consumption and higher performance per watt.....	12
Higher Quality Gaming Experience .....	14
Highly Responsive User Interfaces and Faster Multitasking .....	18
Conclusion.....	21
Document Revision History.....	22

## **Executive Summary**

Desktop CPU manufacturers transitioned to multi-core processor architectures five years ago to address the growing performance demands and the exponential growth of power consumption of single core processors. By using multiple cores the CPUs of today can complete more work faster, and at lower power, than their single core predecessors.

Mobile processors are facing the same performance and power challenges. People use their mobile device much like they use their PC, and expect a similar level of capabilities – while maintaining and increasing battery life. Mobile use cases such as HD video playback, streaming video and audio, multitasking, browsing the web, 3D gaming, and 3D interfaces are stretching the capabilities of current single core mobile processors.

In order to further increase the performance, and extend battery life, mobile devices will transition to multi-core CPUs.

## **The Need for Multiprocessing**

Mobile devices perform a wide variety of tasks such as Web browsing, video playback, mobile gaming, SMS text messaging, and location-based services. Due to the growth in the availability of high speed mobile and Wi-Fi networks, mobile devices will also be used for various performance-intensive tasks that were previously handled by traditional PCs. The next generation of smartphones (called “Super phones”) and tablets will be used for a wide variety of tasks such as playback of high definition 1080p videos, Adobe® Flash®-based online gaming, Flash-based streaming high definition videos, visually rich gaming, video editing, simultaneous HD video downloads, encode and uploads, and real-time HD video conferencing.

The current generation of mobile processors is not designed to deal with this tidal wave of high performance use cases. The quality of experience on devices based on single core CPUs rapidly degrades when users run several applications concurrently, or run performance-intensive applications such as games, video conferencing, video editing, and more. In order to improve CPU performance, engineers employ several techniques, such as using faster and smaller semiconductor processes, increasing core operating frequency and voltage, using larger cores, and using larger on-die caches.

Increasing the size of the CPU core or cache delivers performance increases only up to a certain level, beyond which thermal and heat dissipation issues make any further increase in core and cache size impractical. From basic semiconductor physics we know that increasing operating frequency and voltage can exponentially increase power consumption of semiconductor devices. Even though engineers may be able to squeeze out higher performance by increasing frequency and voltage, the performance increase would drastically reduce battery life. In addition, processors that consume higher power would require larger cooling solutions resulting in an undesired expansion in device size. Therefore, increasing the operating

frequency of the processor to meet the ever-increasing performance requirements of mobile applications is not a viable solution for the long run.

To satisfy the rapidly growing demand for performance and form factor sleekness in mobile devices, the industry has begun to adopt newer technologies such as Symmetrical Multiprocessing and Heterogeneous Multi-core computing. NVIDIA Tegra is the world's most advanced mobile processor built from the ground up as a heterogeneous multi-core SoC (System-On-a-Chip) architecture with two ARM Cortex A9 CPU cores and several other purpose-built cores to handle specialized tasks such as audio, video, and graphics. Purpose-built cores require fewer transistors, operate at lower frequencies, deliver higher performance, and consume much lower power than general purpose processing cores for tasks such as audio, video, and graphics processing.

## **Symmetrical Multiprocessing (SMP)**

Symmetrical Multiprocessing technology enables mobile processors to not only deliver higher performance, but also meet peak performance demands while staying within mobile power budgets. A multi-core architecture with SMP is defined by the following characteristics:

- Architecture consists of two or more identical CPU cores.
- All cores share a common system memory and are controlled by a single Operating system.
- Each CPU is capable of operating independently on different workloads and whenever possible, is also capable of sharing workloads with the other CPU.

Imagine a mobile phone that has a dual core CPU with SMP support-- if the phone's navigation application is running concurrently with a streaming audio application, the OS can assign the navigation task to one CPU core and the streaming audio task to the second CPU. Another example is a single multi-threaded application that can benefit from multiple CPUs. The OS can assign the threads to run on both CPUs concurrently and finish the task faster by sharing the workload across the two CPUs. Since the workload is split across the two cores, these cores can run at a reduced speed while achieving excellent performance and also conserving power (running at lower frequency lessens the voltage required, resulting in a reduction in power by the square of voltage decrease).

Figure 2 below illustrates a typical webpage format that is widely used on the Internet. It contains several ActiveX and JavaScript-based menu options, and includes two embedded Flash videos and Flash animation. A single core CPU would have to bear the entire load of processing the ActiveX/JavaScript content, Flash player processing, and the embedded video decode tasks. In addition, the processor will have to handle background tasks such as Twitter® streams, voice applications, navigation, and others. Under these heavy multitasking conditions the processor is typically running at maximum frequency and voltage, therefore consuming large amounts of power.

On a mobile processor that has a dual core CPU, the tasks to process the various elements of the webpage can be shared between the two processing cores. For example, one core could process background system tasks and Flash video content while the other core could handle ActiveX content, Flash Animation, and video. Figure 3 below illustrates how the various processing tasks could be shared across the two cores. Due to task sharing, the cores don't need to run at full capacity and can be run at a lower frequency and voltage. Since the power consumption of semiconductor devices is proportional to the frequency and voltage-squared, even a small reduction in the operating frequency and voltage will result in significant reduction in power consumption. Therefore a mobile processor with a dual core CPU with SMP capabilities will often be more power efficient than a single core CPU based mobile processor.

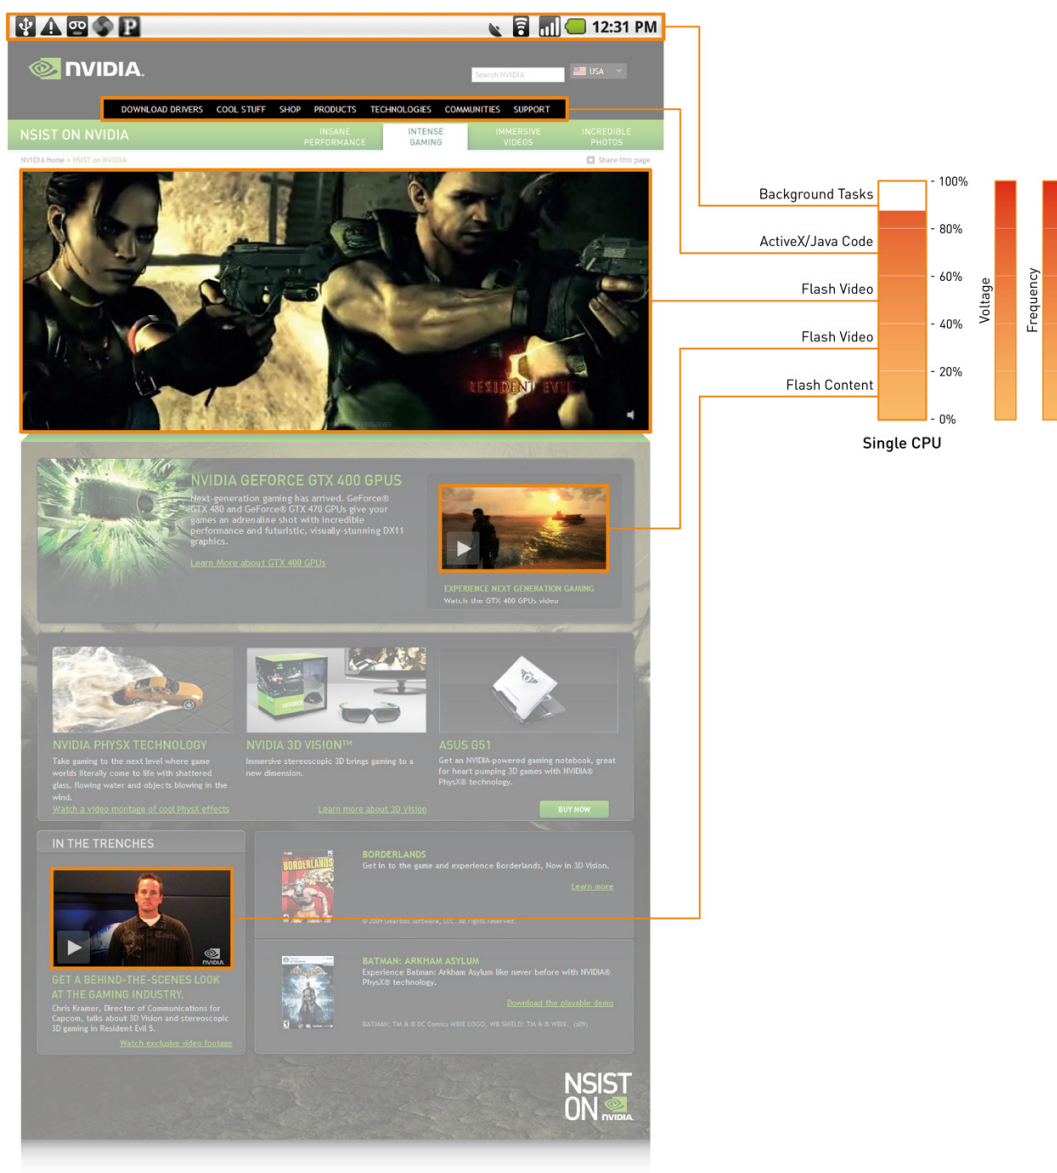


Figure 1 - Typical CPU utilization on single core mobile devices for Web browsing

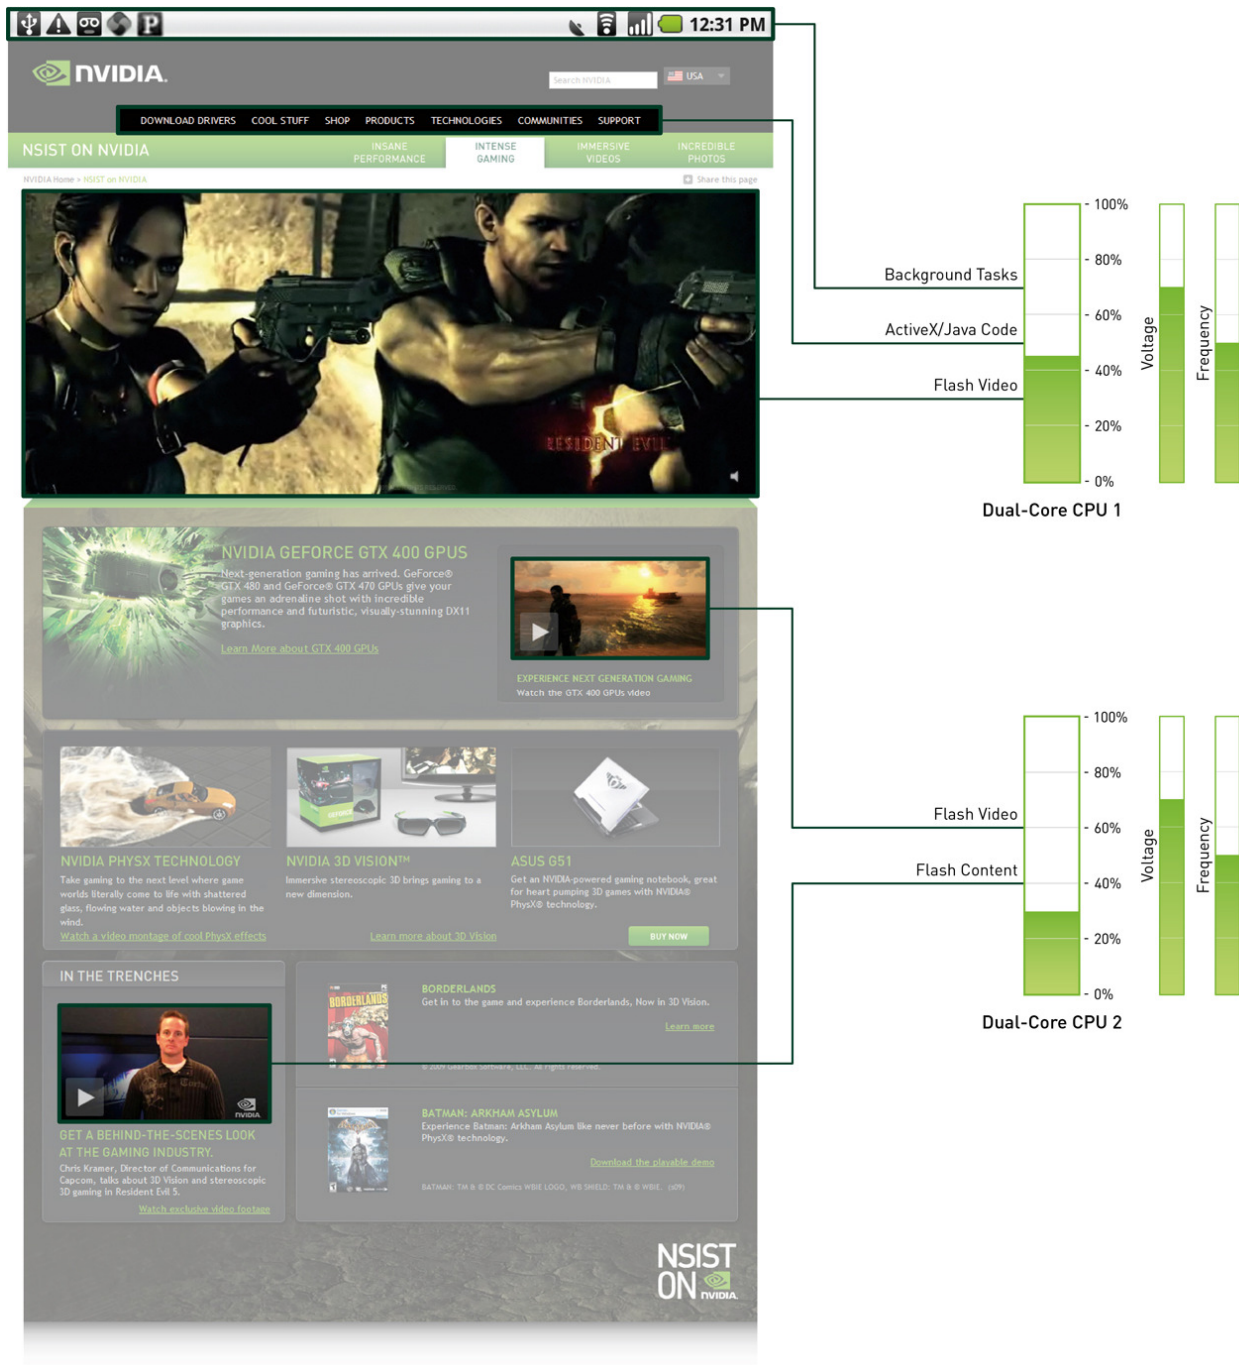


Figure 2 - CPU utilization on a dual core CPU based mobile processor for Web browsing

## **Introducing NVIDIA Tegra 2: The world's first mobile processor with a Dual-Core CPU**

NVIDIA Tegra is the world's first mobile processor that includes a dual core CPU. The dual core CPU in NVIDIA Tegra is a highly optimized version of the ARM® Cortex A9 MPCore™ architecture that delivers almost two times the performance of current mobile processors. The Symmetric Multiprocessing, out of order execution, and superior branch prediction features of the optimized ARM cores help deliver very fast Web page load times, snappy webpage rendering, and a silky smooth user interaction experience.

### **Dual-core ARM Cortex A9 Architecture**

The Cortex-A9 MPCore processor implements the ARMv7 instruction set architecture and is designed around an advanced and highly efficient out-of-order eight-stage pipeline. The Cortex-A9 MPCore processor delivers higher performance over previous generation ARM CPUs and at the same time remains within the power budgets required for mobile devices. The ARM Cortex A9 architecture has evolved from the previous generation ARM Cortex A8™ processor and includes several performance enhancing features such as:

- Dynamic length 8-stage pipeline supporting speculative out-of-order execution. This allows the processor to dynamically reorder instructions to improve performance by avoiding stalls due to instruction latencies and resource conflicts. Older generation Cortex-A8 processors use an in-order pipeline and are unable to avoid the penalties that arise from branching and cache misses
- Support for speculative branch predictions to avoid branching penalties.
- A Dual-core Symmetrical Multiprocessing (SMP) configuration operating either independently, or in lockstep to deliver peak performance when needed, and consuming almost zero power when idle.
- 32KB Instruction cache and 32KB Data cache per core with both cores sharing a common 1MB L2 Cache. The 1MB L2 cache is large enough to load an entire browser memory footprint into cache to provide a faster Web browsing experience.
- CPU cores that are optimized to operate at a frequency of one Gigahertz with the ability to scale up to even higher frequencies. The two cores are assisted by a common snoop control unit that enforces coherency between the cores and manages the common 1MB L2 cache shared by the two cores.

Both CPU cores are power managed through complex and highly intelligent Dynamic Voltage and Frequency Scaling algorithms. These algorithms are implemented at both the hardware and software level to ensure both cores are always operating at the optimal voltage and frequency levels to deliver the demanded performance, while consuming the lowest possible power. The algorithms were developed by NVIDIA after extensive study of mobile use cases to ensure that

the cores always operate at maximum efficiency, and are not only able to meet peak performance needs, but also stay in ultra low power mode when the system is idle.

A block diagram of the dual-core Cortex A9 microprocessor is shown in Figure 1.

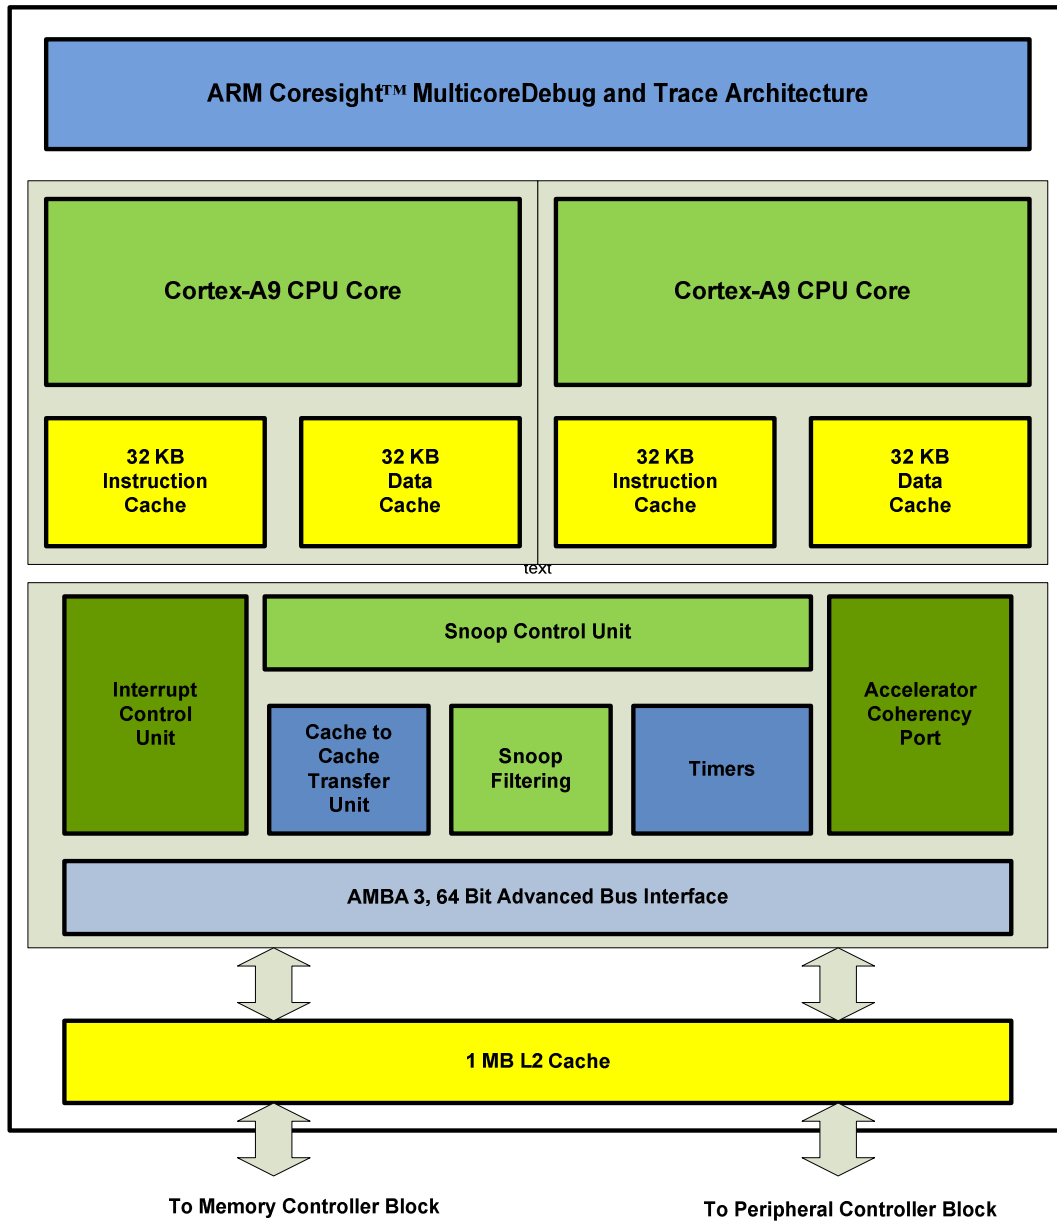


Figure 3 - Dual Core ARM Cortex A9 MPcore Implementation




## Benefits of Symmetrical Multiprocessing in Tegra 2

Symmetrical Multiprocessing in Tegra 2 delivers not only higher performance and lower power consumption, but also tangible benefits to several mobile use cases. Some of the benefits are:

- Faster Web page load times
- Lower power consumption and higher performance per watt
- Higher quality game play experience for advanced console-style mobile games
- Highly responsive and smoother UIs (user interfaces)
- Faster multitasking

### *Faster Web page load times*

Mobile devices that have a multi-core CPU with SMP support will be able to deliver an ultra fast desktop PC style Web browsing experience. Gone are the days when browsers run as a single process with a single thread. Most modern browsers such as Google Chrome® and Mozilla Firefox® are now multi-threaded and are capable of spawning several processes concurrently. Each page tab in a Chrome browser is a separate process, and each process manages its own set of threads. Both the processes and threads are highly parallelized.



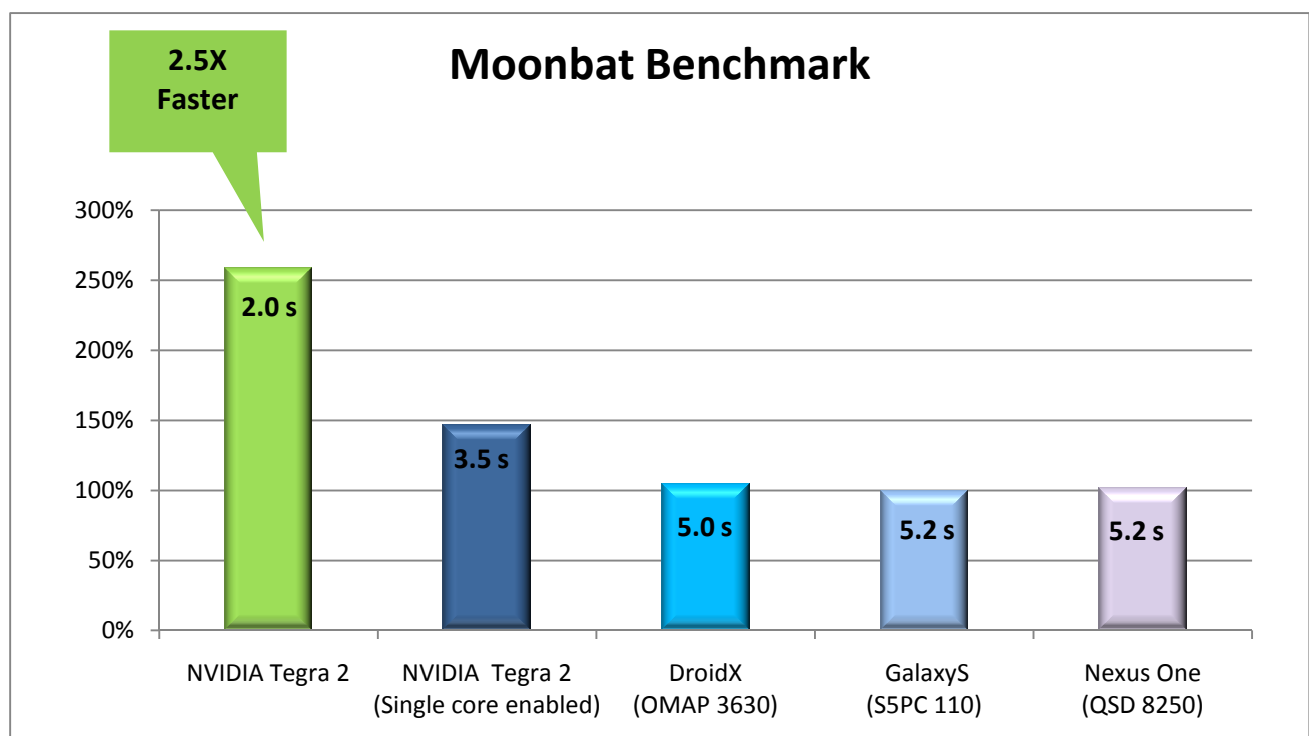
**Figure 4 - Multi-threaded browsers allocate tasks across both cores**

NVIDIA Tegra 2 with its dual core CPU and SMP technology is able to run these processes independently on separate CPU cores to deliver Web page rendering and page load times that are significantly faster than the competition.

Results from various widely accepted browser and JavaScript benchmarks confirm that NVIDIA Tegra 2 delivers significantly faster runtimes and page render times than competing single core CPU based solutions. Results from Moonbat, a multi-threaded JavaScript benchmark show that the dual core Cortex A9 based Tegra delivers almost 1.7x faster runtime than a single core Cortex A9 based Tegra processor and almost 2.5x faster than competing single core CPU based mobile processors.. Results from Coremark, a widely accepted CPU performance benchmark confirm that the dual core Cortex A9 in the NVIDIA Tegra delivers twice the performance of a single core cortex A9 processor.

The performance improvement provided by dual core CPU in NVIDIA Tegra 2 was also publicly demonstrated at the ARM Technology Conference 2010 and the ARM representative in the demo video<sup>1</sup> mentions that Web page load time on the dual core Cortex A9 based Tegra 2 processor is at least 50% faster than a single core Cortex A9 processor

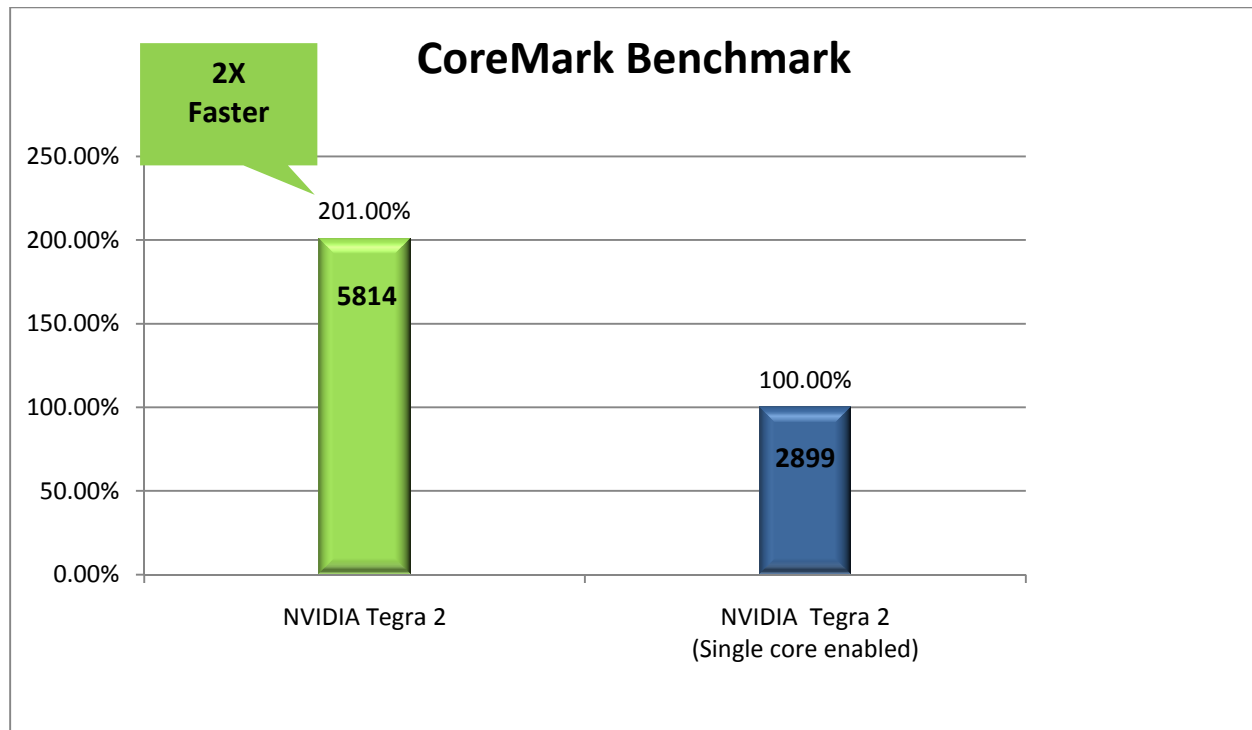
The following set of benchmark results show that the Dual Core Cortex A9 based NVIDIA Tegra processor is 1.5x to 2x faster than an equivalent single core Cortex A9 version and at least 1.7x to 2.4x faster than competing ARM Cortex A8 based application processors.



**Figure 5 - Dual Core A9 Performance on Moonbat JavaScript Benchmark<sup>2</sup> (Lower Score wins)**

<sup>1</sup> <http://armdevices.net/2010/11/13/nvidia-tegra2-arm-cortex-a9-dual-core-performance-for-web-browsing/>

<sup>2</sup> Android OS: 2.2, Web Browser: Firefox



**Figure 6 - NVIDIA Tegra Performance on Coremark CPU Benchmark<sup>3</sup> (Higher Score Wins)**

Due to the growth of Web2.0 technologies, a Web page today no longer contains simple links to resources. Instead, a Web page may have a series of scripts that need to be executed and parsed to obtain links to additional resources that must be downloaded to create the page. Correspondingly, browsers are capable of executing multiple scripts in parallel, and setting up multiple concurrent connections to servers to load many resources in parallel. If a browser running on a device with only one CPU core encounters a script on a page, it must execute that single script and cannot parse the rest of the page, run other scripts on that page, initialize parallel connections or fetch additional resources until it finishes executing the first script. In a Web page with multiple scripts, it leads to a situation where a script is the only resource that is loading, and connection-level parallelism offered by today's browsers is not exploited.

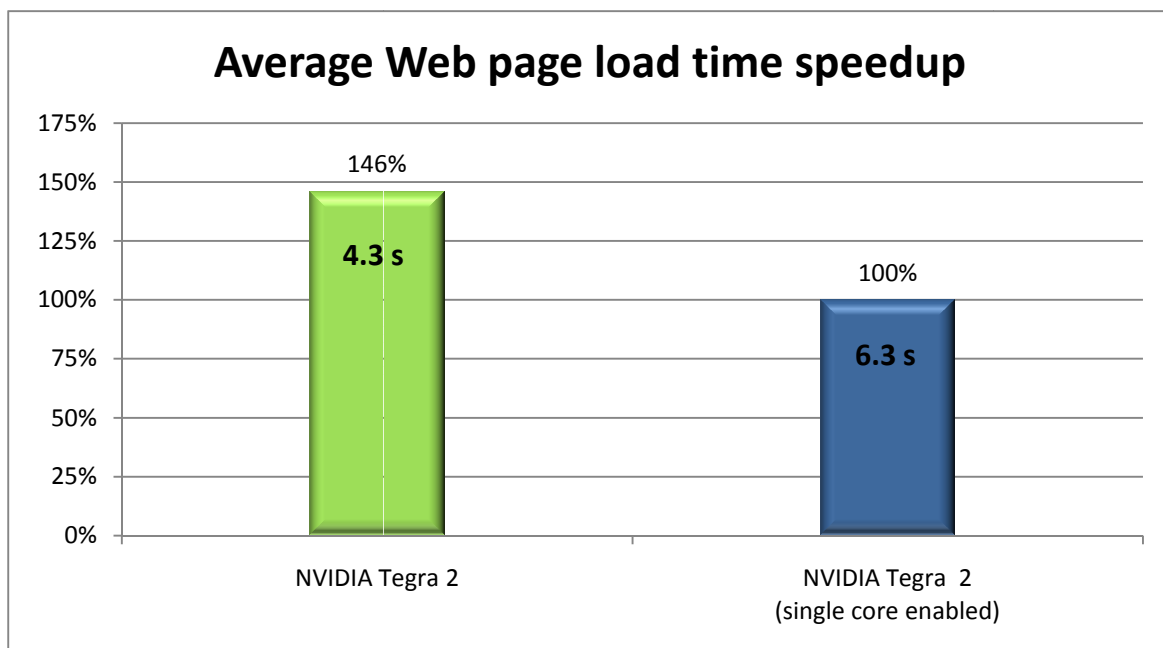
The situation described above can be significantly improved in an SMP environment. A popular open source browser layout engine called Webkit is used by web browsers such as Chrome, Safari, and Android Web Browser. Webkit provides a set of classes to display Web content in windows and implements various browser features. Recent optimizations in the Webkit code allow SMP capable processors such as NVIDIA Tegra to launch multiple parsers in parallel on

<sup>3</sup> Dual Core: CoreMark 1.0 : 5814.305896 / GCC4.4.1 (Sourcery G++ Lite 2010q1-202) -O3 -mcpu=cortex-a8 -funroll-loops -falign-loops=8 -fgcse-sm -fno-tree-vectorize -marm / Heap / 2:PThreads

Single Core: CoreMark 1.0 : 2899.759693 / GCC4.4.1 (Sourcery G++ Lite 2010q1-202) -O3 -mcpu=cortex-a8 -funroll-loops -falign-loops=8 -fgcse-sm -fno-tree-vectorize -marm / Heap / 2:PThreads

the two processing cores when the main parser encounters a script on the Web page. These mini parsers speculatively parse the rest of the page and go fetch the rest of the resources from the server in parallel.

On a non-SMP processor, this would not be possible because everything has to be processed sequentially, and the mini parsers will not be able to execute until the main script finishes its task. Thus mobile processors with SMP capabilities will be able to deliver much faster page load times. The figure below shows the average Web page load time measured across popular websites and it can be seen that the page load time on a Dual Core Cortex A9 based Tegra device is fifty percent faster than on an equivalent Single Core Cortex A9 based Tegra device.



**Figure 7 - Average Speed Up in Web Page load time on a Dual Core Tegra Processor<sup>4</sup>**

NVIDIA Tegra is the only mobile processor in the market today that has a dual core processor with SMP technology and it delivers the fastest and richest Web browsing experience on a mobile device.

#### ***Lower Power consumption and higher performance per watt***

A common misconception is that the NVIDIA Tegra dual core CPU architecture consumes more power than competing single core CPU solutions and causes significant reduction in battery life. On the contrary, due to SMP and intelligent power management algorithms, the Tegra solution is more power efficient and delivers higher performance per watt than single core processors. In order to meet peak performance demands in a multitasking environment, a single core CPU not only runs at higher clock frequencies and voltages than a dual core CPU, but also takes longer periods of time to complete a given task. The Cortex A9 dual core CPU implementation in

<sup>4</sup> Averaged over 10 samples at 800x480 resolution on an offline web server using Webkit browser and Android 2.2

NVIDIA Tegra is more power efficient and delivers higher performance per watt than competing solutions for the following reasons:

- NVIDIA Tegra employs SMP technology to distribute and share task workloads across the two processing cores and thus each core is not fully loaded and does not have to run at peak capacity/speed. This enables the system level power management control logic to run the two cores at much lower operating frequency and voltage and thus achieve significant power savings
- For tasks that are highly parallel, NVIDIA Tegra is able to distribute the workload across the two CPU cores and complete the task much faster than a single core CPU solution. Thus the dual core CPU on NVIDIA Tegra would be able to complete a task quickly and enter into a low power state to conserve power, while a single core processor would have to be in an active high power state for longer periods of time to process the same task.
- For low intensity workloads that only require the processing power of a single core, the other core can be turned off, reducing power consumption to almost the same level as that of a single core CPU.

Consider browsing the Web on a single core CPU. If the Web page contains several scripts, streaming Flash video content, and script-based images, then in most cases, the single core CPU will be running at peak utilization, and to deliver peak performance, it will also be running at maximum operating voltage and frequency. Let us assume that the single core CPU for this task consumes power equal to '**P**'. Let us also assume that the operating voltage of the CPU is **1.1V** and the operating frequency is **1 GHz**.

Now consider the same task running on a dual core Cortex A9 CPU used in the NVIDIA Tegra architecture. Due to Symmetrical Multiprocessing, the Web browsing task is shared between the two Cortex A9 cores. Therefore both cores need to run at only around 50% utilization to complete the task. Since the workload is shared, the cores can run at much lower voltage and frequency. Because each core processes only half the workload, each core can operate at almost half the frequency of the single core CPU, and therefore can run at a lower operating voltage. Let us assume that each core runs at a frequency of **550 MHz** and at an operating voltage of **0.8V** to process the task.

The dynamic power consumption of each core is proportional to the operating frequency and to the square of the operating voltage. Due to the exponential relationship between power and voltage, the dual core CPU consumes lower power than a single core CPU for the same workload. The figure below shows how a dual core Cortex-A9 CPU will consume only 60% of the power consumed by a single core Cortex-A8 for the same task.

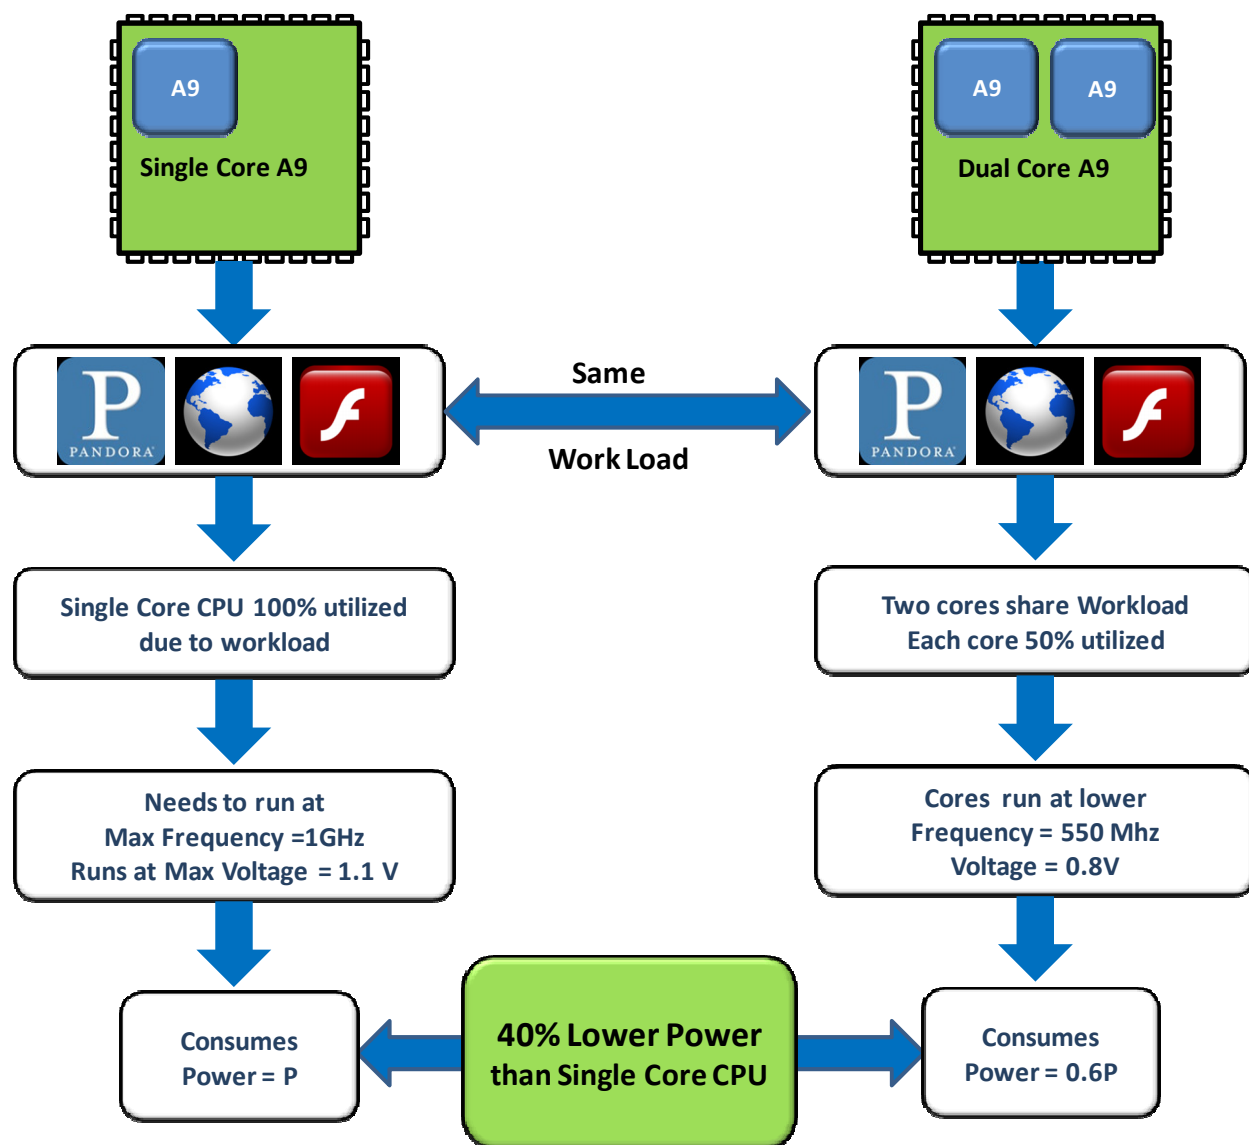


Figure 8 - Voltage and Frequency Scaling Benefits of Dual Core CPU

### ***Higher Quality Gaming Experience***

Web browsing is not the only use case that benefits from a multi-core SMP architecture. Mobile Gaming is a rapidly growing use case, and recent market study data showed that the number of users playing games on smartphones has increased by sixty percent in the 2009-2010 time period.<sup>5</sup> In addition, games played on mobile devices today have evolved from simple 2D versions, and now employ complex 3D rendering techniques. Many visually rich and compelling games from the PC and console platforms are becoming available for mobile devices.

<sup>5</sup> [www.comscore.com/Press\\_Events/Press\\_Releases/2010/4/Smartphone\\_Adoption\\_Shifting\\_Dynamics\\_of\\_U.S.\\_Mobile\\_Gaming\\_Market](http://www.comscore.com/Press_Events/Press_Releases/2010/4/Smartphone_Adoption_Shifting_Dynamics_of_U.S._Mobile_Gaming_Market)

Consoles and PCs today are all multi-core chips that support SMP technology and most console and PC games are coded to take advantage of the multi-core and SMP capabilities of the hardware. Today, most game engines are multithreaded, and the engines are increasingly migrating to task-processing models where the “size” of the individual jobs is reduced while the number of threads is increased. All modern game platforms support many threads and this trend will continue to increase on future platforms.

Due to the recent growth in mobile gaming and the tremendous growth potential for gaming on mobile devices, PC and console game developers are now developing or porting over popular console and PC games to mobile devices. Most of the PC and console games are optimized for multi-core hardware and therefore will deliver the best gaming experience on multi-core and SMP capable mobile devices.

*“Because Tegra is a dual core CPU, we can run our 2500 level AI in the background on the other core, all with complete polish and complete smoothness.”*

**- Thomas Williamson, Wardrum Software**

NVIDIA Tegra is the only mobile processor that not only has a dual core SMP capable CPU but also a console class Ultra Low Power (ULP) GeForce™ GPU. The architecture of the GPU in the Tegra processor is similar to that of desktop GeForce GPUs and therefore games that were originally developed for SMP-capable desktop CPUs and desktop GPU architectures can be easily ported to run on NVIDIA Tegra and will deliver the best gameplay experience on Tegra based mobile devices. Table 1 shows some of the popular game engines and the number of threads used by each engine.

Game/Engine	Number of Threads
Unreal Engine 3	4+
Id Tech 5	6+
Frostbite	14
Civilization 5	12
Mafia 2	4
Crysis	8
Uncharted 2	8
Killzone 2	8+

**Table 1 - Multi-threaded Games and Engines**

*Unreal Engine 3.0* is a widely used multi-threaded engine and it uses individual threads for functions such as Render, Audio, Transcode, Collision, Transparency, etc. A SMP capable OS



can distribute these threads across multiple CPUs, remove performance bottlenecks and deliver much smoother game play experiences. Experiments run on games based on the Unreal Engine 3.0 show that when the additional CPU core and SMP is enabled on a Tegra processor, the game runs almost seventy percent faster.



Figure 9 - Dungeon Defender using both cores of the NVIDIA Tegra Processor

*Id Tech 5* is a multi-threaded game engine that is used in games such as *Doom 4* and *Rage*. The engine employs two conventional threads for the primary game processing and Rendering functionalities. It also employs several asynchronous threads for access to the Network, disk drives and Audio. All other functions are coded as individual "jobs" that can be run in parallel on multiple cores. The engine is architected to scale to any number of cores and can easily add or remove "job" classes based on the number of available cores. The engine has defined job classes for tasks such as Collision detection, Animation blend, Obstacle avoidance, Virtual Texturing, Model detail generation, etc.

*Frostbite* is another example of a game engine that employs job-based parallelism. This engine is used by the popular *Battlefield: Bad Company* series of games. It is an engine that is capable of using as many threads as the underlying hardware platform provides. The engine performs the primary Game and Render tasks on the GPU and divides up the other system related work into jobs. Each job typically consists of 15K to 200K lines of C++ code with the average job size being around 25K lines of code. Most of these jobs are independent while some have




interdependencies. Each frame of the game would typically contain two hundred to three hundred jobs and the engine assigns the jobs to all available hardware cores.



**Figure 10 - Frostbite Multithreaded Game Engine employing job level parallelism**

The figure below shows the job distribution across various cores captured during real-time game play. Each color in the figure represents a job launched by the Frostbite Engine. This data was collected using Frostbite Timing View software on a PC system that had a quad core CPU and two GPUs in AFR mode. From the figure it is easily seen that the game engine effectively uses all available cores for task processing. Thus these game engines will be able to efficiently utilize the two CPU cores and the ULP GeForce GPU core of the NVIDIA Tegra and deliver an uncompromised console style gaming experience.



**Figure 11 - Task Level Parallelism on Frostbite Game Engine**

Gaming is a rapidly growing use case on mobile devices such as smartphones, tablets, and portable media players. It is estimated that end-user spending on mobile gaming will grow from \$5.6 billion in 2010 to \$11.4 billion worldwide in 2014<sup>6</sup>. The number of mobile gamers in the United States alone is estimated to grow by almost 50% between 2010 and 2014, reaching 100 million by 2014<sup>7</sup>. The widespread acceptance of smartphones and tablets with increased processing capabilities, higher resolution displays, larger screens, and faster connectivity will further fuel the growing interest and acceptance of mobile gaming. Developers of popular console and PC games will not be able to ignore this rapidly growing market and will begin porting popular console and PC games to mobile phones and tablets.

Most console and PC games were developed for multi-core hardware systems and therefore these games will deliver the best performance on mobile processors that support multi-core architectures and SMP. In addition, developers can more easily port multi-threaded game engines to mobile processors that include multi-core CPUs and GPUs that are similar to PC and console GPUs. NVIDIA Tegra mobile processor includes both a dual core SMP capable CPU and an ULP GeForce class GPU. Therefore, mobile devices based on NVIDIA Tegra will not only deliver the best game play experience, but also will offer the widest variety of console and PC style games.

### ***Highly Responsive User Interfaces and Faster Multitasking***

Another important and highly visible benefit from the SMP capability of a dual core CPU is that the responsiveness of the device to user interaction is significantly improved. User interaction with the mobile device through menus, buttons, and touch events is highly sensitive to latency. Any delay in processing and responding to user input would result in response lag that makes the device interaction an annoying experience.

<sup>6</sup> <http://www.gartner.com/it/page.jsp?id=1370213>

<sup>7</sup> <http://www.mobilemarketer.com/cms/news/research/7096.html>

A variety of factors contribute to poor responsiveness. For example, a fully utilized CPU, unavailability of memory bandwidth, non-optimized driver stacks, etc., all could cause stalling of response to user inputs. Smartphone users typically have several applications running concurrently. For example, it is not uncommon to see applications such as Web browsers, streaming music, email syncs, social network syncs, and news feeders running concurrently on a mobile device. Under such heavy multitasking conditions, single core CPUs often hit peak utilization and are unable to immediately switch over to processing user interaction tasks, and this results in delays and noticeable lag in responsiveness.

An OS with SMP can dynamically allocate workloads to the appropriate CPU core based on current CPU loads and task/activity priority. Thus in a dual core SMP environment, the OS can assign one CPU core to handle tasks such as streaming music and background syncs, and free up the other CPU core to handle latency-sensitive tasks such as user interactions and Web browsing.

Consider a typical use case where a user is playing a graphics intensive game and at the same time listening to music streamed via a popular application such as Pandora® and also is downloading content from the Internet. On a single core CPU, all three tasks would need the immediate attention and processing power of the CPU, which would slow down or interfere with the smooth operation of the tasks. Gaming is a performance-intensive use case that requires a significant portion of the CPU processing power. But due to the addition of real-time tasks such as streaming music and file download that also require CPU processing, a single core CPU will not be able to deliver the performance demanded by the gaming application.

On a dual core CPU with SMP capability, the Operating System will be able to allocate one CPU core for the gaming task and the other CPU core for the streaming music and file download tasks. Thus the gaming task will get the entire processing power of one core and can deliver a superior gaming experience. The figures below illustrates the FPS (Frames Per Second) improvement that SMP and the additional core delivers for the games Quake and Dungeon Defender under the above-described multitasking use case. Because the music playback and file download tasks are offloaded to the second core, the dual core Cortex A9 based Tegra is able to deliver 60% to 100% higher FPS than an equivalent single core A9 based Tegra.




Figure 12 - Quake game play FPS improvement due to dual core A9 under heaving multitasking conditions<sup>8</sup>




Figure 13 - Dungeon Defender Game Play FPS improvement due to dual core A9 under heavy multitasking<sup>9</sup>

<sup>8</sup> At 800x480 res, Android 2.2, Pandora music live streaming, file download via USB

<sup>9</sup> At 800x480 res, Android 2.2, Pandora music live streaming, file download via USB

## Conclusion

For several decades PC processors were primarily single core architectures and CPU makers increased performance by increasing operating frequencies, core sizes, and using smaller manufacturing processes allowing more transistors in the same chip area. However, PC manufacturers realized that the continual increase in frequency and core sizes caused exponential increases in power consumption and excessive heat dissipation. Therefore, CPU makers developed multi-core CPU architectures to continue delivering higher performance processors, while limiting the power consumption of these processors. Most desktop and notebook systems today use either a dual or quad core processor and consume significantly lower power than their single core predecessors. But mobile devices like smartphones and tablets benefit even more from multi-core architectures because the battery life benefits are so substantial. Dual-core processors will be the standard in 2011, and quad-core is coming in the near future.

Mobile application processors are facing the same performance and power challenges that the desktop and notebook CPUs faced a few years ago. Demanding mobile applications such as HD video playback, streaming videos, 3D gaming, 3D interfaces, etc., are already stretching the capabilities of current single core mobile processors. In order to further increase the performance and stay within mobile power budgets, it is inevitable that all mobile processors will eventually have multi-core processors. Mobile operating systems such as Android, Windows® CE, and Symbian are capable of operating in a multi-core environment, and have the features required to efficiently harness the multiple processing cores of the underlying hardware. Also, popular Web browsers and most PC games are already multithreaded, and users will see tremendous improvements in performance if these applications are ported to run on multi-core CPU-based mobile processors.

NVIDIA Tegra was designed to harness the power of Symmetrical Multiprocessing and deliver a phenomenal Web browsing experience, highly responsive user interfaces, effective multitasking, and tremendous battery life improvements.

## **Document Revision History**

**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, CUDA, FERMI and GeForce 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**

© 2010 NVIDIA Corporation. All rights reserved.