Whitepaper

Variable SMP (4-PLUS-1™)

– A Multi-Core CPU Architecture for Low Power and High Performance

Table of Contents

Introduction

In February 2011, NVIDIA introduced and demonstrated its Tegra 3 mobile processor, the world's only 4-PLUS-1™ quad core mobile processor. Tegra 3 will enable new mobile applications, new experiences, more robust multitasking, higher-quality gaming, and faster web browsing. In addition, Tegra 3 will further extend battery life by operating its CPU cores at lower frequency, yet still be able to complete more work than dual core or single core processors.

Other industry leaders have agreed that quad core is the correct direction, and announced their own mobile quad core processor lines. While including more CPU cores will improve performance and lower power in many use cases, additional technologies can take this even further.

NVIDIA's Tegra 3 processor implements a novel new **Variable Symmetric Multiprocessing (vSMP)** technology. Not previously disclosed publicly, vSMP includes a fifth CPU core (the "Battery Saver core) built using a special low power silicon process that executes tasks at low frequency for active standby mode, music playback, and even video playback. The four main "quad" cores are built using a standard silicon process to reach higher frequencies, while consuming lower power than dual core solutions for many tasks. All five CPU cores are identical ARM Cortex A9 CPUs, and are individually enabled and disabled (via aggressive power gating) based on the work load. The "Battery Saver" core is OS transparent, unlike current Asynchronous SMP architectures, meaning the OS and applications are not aware of this core, but automatically take advantage of it. This strategy saves significant software efforts and new coding requirements.

Optimized for Key Mobile Use Cases

Mobile use case studies show that most mobile devices are typically in active standby state for eighty percent of the time, and process intensive mobile applications twenty percent of the time.

Think about your device sitting in your pocket or on a desk as "Active Standby" – or when the user is not actively interacting with the device, and the processor is either running background tasks or low performance applications that do not require user interactions. On the other hand, when you're using your device – browsing the web, checking Email, gaming, running multimedia applications, playing back media, and so on – the device is in a performance-intensive mode that may require one or more CPU cores running at higher frequency ranges.

Keep in mind, when the device is in active standby state many tasks are still happening in the background - Email syncs, social media syncs, live wallpapers, active widgets, and more. Such tasks can be processed by just a single CPU core running at much lower frequencies. Users generally do not care how fast the background tasks are processed, only that they happen and do not consume much battery life.

When in active standby state, battery life can be improved significantly by minimizing the active standby state power consumption of mobile processors.

Figure 1 Background tasks that are typically running when device is in active standby state

Silicon Process and its Impact on Power and Frequency

Power consumption of a silicon device is equal to the sum of leakage power and dynamic power. Leakage power is mainly determined by the silicon process technology, and dynamic power is determined by silicon process technology and by operating voltages and frequencies.

Dynamic power of a silicon device is proportional to the operating frequency, and more importantly proportional to the square of the operating voltage.

Total Power = Leakage Power + Dynamic Power

Dynamic Power α Frequency x Voltage^2

When a silicon device is operating at or near peak frequency, the total power consumption of the device is dominated by the dynamic power consumption, and when the device is idle or operating at near idle conditions, the leakage power is a significant portion of the total power consumption.

Transistors on fast process technologies consume high leakage power and have very fast switching times at normal voltage levels. Therefore, CPU cores built on fast process technologies (CPU A in Figure 2) consume high leakage power under idle or active standby

conditions, but are capable of running at higher frequency ranges without requiring significant increases in operating voltage.

Transistors on low power process technologies have low leakage power but have slower switching times at normal voltage levels, and to enable them to switch faster (for high frequency operations) they need higher than normal voltage ranges.

CPU cores built on low power process technologies (CPU B in Figure 2) consume very low leakage power but require higher than normal voltage levels to operate at very high frequencies. Therefore, they consume excessive amounts of dynamic power and can cause significant power and thermal issues.

The following simplified statements effectively convey the concept:

Fast Process = Optimized for high frequency operation, but higher leakage

Low Power Process = Operates at lower frequency with lower leakage

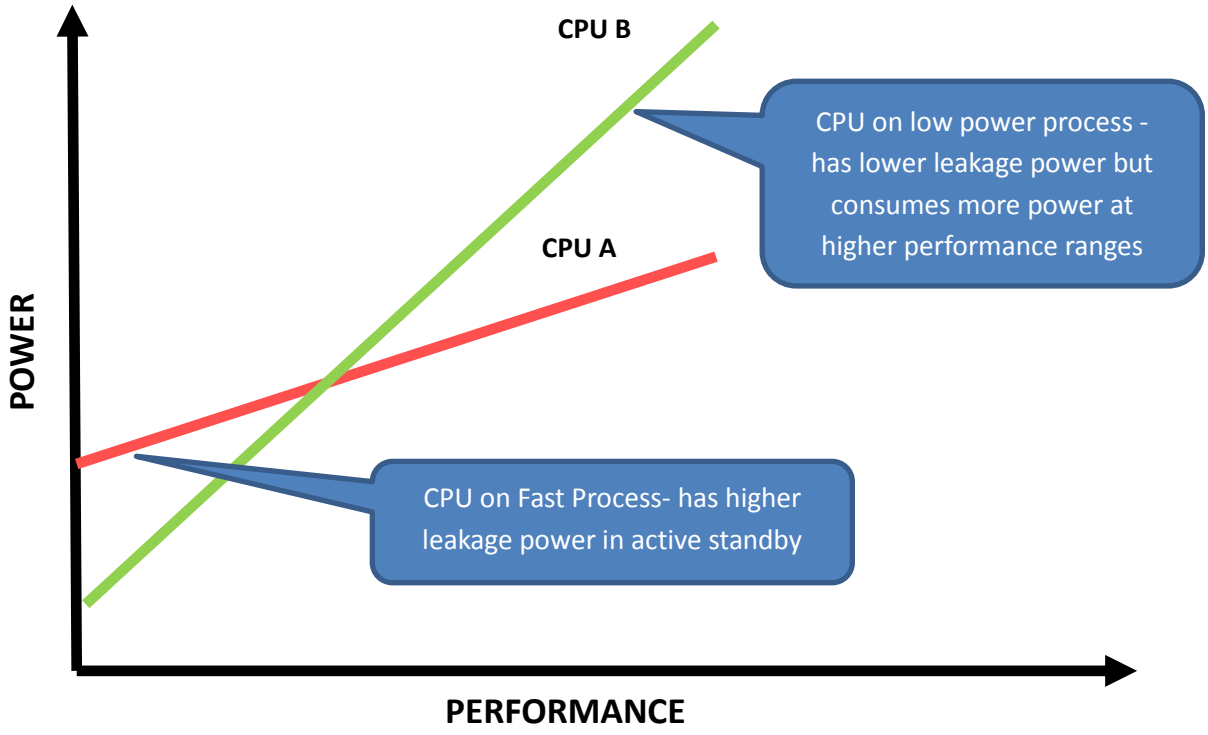

Figure 2 Mobile CPU Power-Performance curves

To meet the rapidly increasing demands of performance-intensive mobile use cases and extended battery life, it is becoming increasingly difficult to minimize both the active standby and dynamic power consumption of CPU cores. By combining both silicon processes (defined

above), along with architectural optimization, a single SoC (System-on-a-Chip) can be optimized for both high performance and low power consumption.

Variable Symmetric Multiprocessing – 4-PLUS-1 quad core architecture

NVIDIA's Tegra 3 is the world's only mobile SoC device to implement a patented **Variable Symmetric Multiprocessing (vSMP)** technology that not only minimizes active standby state power consumption, but also delivers on-demand maximum quad core performance, we call this our 4-PLUS-1 quad core architecture. In addition to four main Cortex A9 high-performance CPU cores, Tegra 3 has a fifth low power, low leakage Cortex A9 CPU core called the '**Battery Saver**' CPU core that is optimized to minimize active standby state power consumption, and handle less demanding processing tasks.

NVIDIA Tegra 3 also includes other patented vSMP technologies that intelligently manage workload distribution between the main cores and the Battery Saver core based on application and operating system requirements. This management is handled by NVIDIA's Dynamic Voltage and Frequency Scaling (DVFS) and CPU Hot-Plug management software and does not require any other special modifications to the operating system.

Low Power Battery Saver Core

The Battery Saver core is designed on a low power process technology, but has an identical internal architecture as the main Cortex A9 CPU cores. Since it is built on a low power process in the low performance ranges (and frequencies), it consumes lower power than the main CPU cores that are built on a fast process technology. Power-performance measurements on Tegra 3 show that the Battery Saver core delivers higher performance per watt than the main cores at operating frequencies below 500 MHz, and therefore the maximum operating frequency of the Battery Saver core is capped at 500MHz. Table 1 compares and contrasts the Battery Saver core to the four main cores on Tegra 3.

	Power optimized Battery Saver CPU Core	**Performance optimized main CPU Cores**
Architecture	Cortex A9	Cortex A9
Process Technology	Low Power (LP)	General/Fast (G).
Operating Frequency Range	0 MHz to 500 MHz	0 MHz to Max GHz

Table 1 Battery Saver and Main CPU Core features

The Battery Saver core is used primarily when the mobile device is in active standby and performing background tasks such as Email syncs, Twitter updates, Facebook updates etc. It is also used for applications that do not require significant CPU processing power, such as streaming audio, offline audio, and both online or offline video playback. Note that both audio and video playback, in addition to video encoding, are largely processed by hardware-based encoders and decoders.

Unlike the Battery Saver core, the main CPU cores need to operate at very high frequencies to deliver high performance. Therefore they are built on a fast process technology which allows them to scale up to very high operating frequencies at lower operating voltage ranges. Thus the main cores are able to deliver high performance without significant increases in dynamic power consumption.

Figure 3 Low Power Battery Saver CPU on Tegra 3

Using the combination of performance-optimized main cores and a power-optimized Battery Saver core, Variable Symmetric Multiprocessing (4-PLUS-1) not only delivers ultra-low power consumption in active standby states, but also on-demand peak quad core performance for performance hungry mobile applications such as gaming, Web browsing, Flash media, and video conferencing.

vSMP successfully combines the power-performance benefits of the power-optimized CPU B and performance-optimized CPU A shown in Figure 2 and delivers a power-performance curve that looks like the one shown in Figure 4.

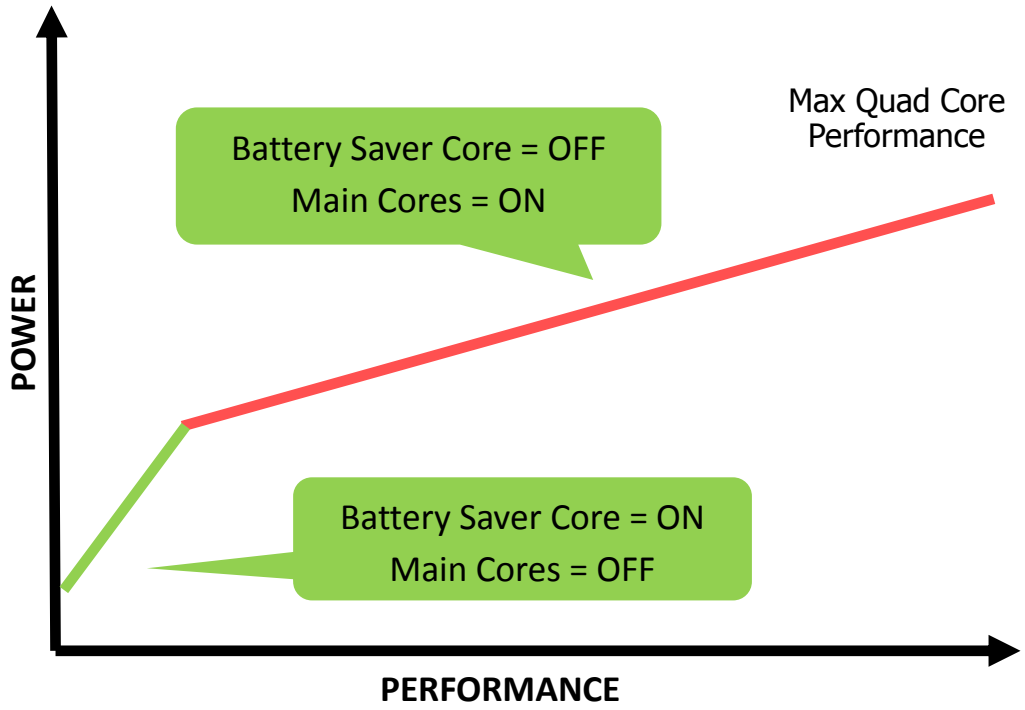

Figure 4 Power-Performance curve of Battery Saver core plus quad main cores running on vSMP technology

Operating System Transparent Implementation

The Android 3.x (Honeycomb) and Android 4.x (Icecream Sandwich) operating system has built-in support for multi-processing and is capable of leveraging the performance of multiple CPU cores. However, the operating system assumes that all available CPU cores are of equal performance capability and schedules tasks to available cores based on this assumption. Therefore, in order to make the management of the Battery Saver core and main cores totally transparent to the operating system, Tegra 3 implements both hardware-based and low level software-based management of the Battery Saver core and the main quad CPU cores.

Patented hardware and software CPU management logic continuously monitors CPU workload to automatically and dynamically enable and disable the Battery Saver core and the main CPU cores. The decision to turn on and off the Battery Saver and main cores is purely based on current CPU workload levels and the resulting CPU operating frequency recommendations made by the CPU frequency control subsystem embedded in the operating system kernel. The technology does not require any application or OS modifications.

Workload-Based Dynamic Enabling and Disabling of CPU Cores

When the Battery Saver core is turned off and the mobile processor is using the main cores for processing, the CPU Governor and CPU management logic continues to monitor CPU workload and utilization of each of the main cores, dynamically enabling or disabling one to four of the main cores. For example, applications such as Email, basic games, or text messaging typically need the processing power of just one of the four main cores. For more demanding applications such as Flash-intensive Web browsing or heavy multitasking, the CPU manager may turn on two CPU cores. But to meet peak performance demands of applications such as console class gaming and media editing or creation, all four CPU cores would be turned on to deliver peak performance demanded by the application(s).

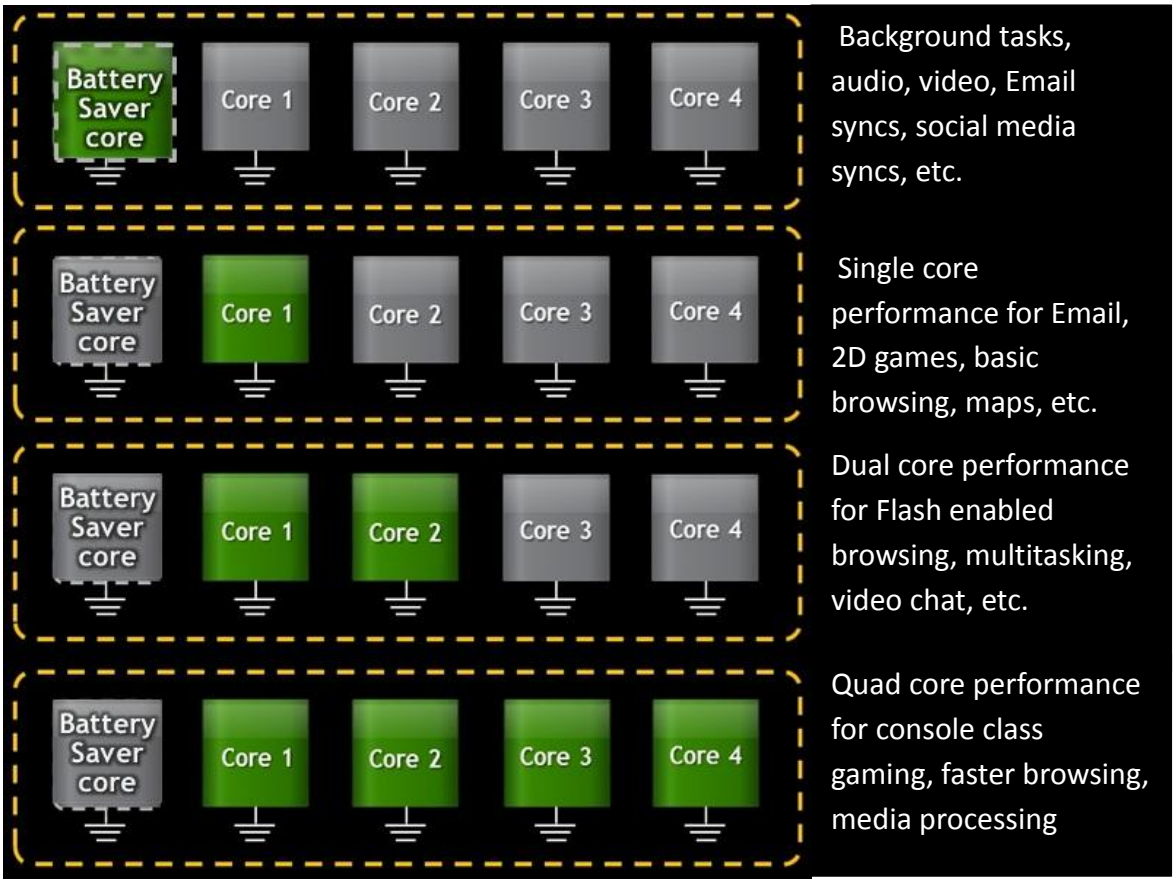

Figure 5 CPU core management based on workload

Architectural Advantages of vSMP Architecture

vSMP technology has several architectural advantages compared to other solutions, such as asynchronous clocking.

- **Cache Coherency:** Since vSMP does not allow both the Battery Saver core and the main cores to be enabled at the same time, there are no penalties involved in synchronizing caches between cores running at different frequencies. The Battery Saver and main cores share the same L2 cache, and the cache is programmed to return data in the same number of nanoseconds for both Battery Saver and main cores (essentially it takes more “main core cycles” versus fewer “Battery Saver core cycles” because the main cores run at higher frequency).
- **OS Efficiency:** The Android OS assumes that all available CPU cores are identical with similar performance capability and schedules workloads to these cores accordingly. When multiple CPU cores are each run at different asynchronous frequencies, it results in the cores having differing performance capabilities. This could lead to OS scheduling inefficiencies. In contrast, vSMP technology always maintains all active cores at a similar synchronous operating frequency for optimized OS scheduling. Even when vSMP switches from the Battery Saver core to one or more of the main CPU cores, the CPU management logic ensures a seamless transition that is not perceptible to end users and does not result in any OS scheduling penalties.
- **Power Optimized:** Each core in an asynchronous clocking based CPU architecture is typically on a different power plane (aka voltage rail or voltage plane) to adjust the voltage of each core based on operating frequency. This could result in increased signal and power line noise across the voltage planes and negatively impact performance. Since each voltage plane may require its own set of voltage regulators, these architectures may not be easily scalable as the number of CPU cores is increased. The additional voltage regulators increase BOM (Bill of Materials) cost and power consumption. If the same voltage rail is used for all cores, then each core will run at the voltage required by the fastest core, thus losing the advantage of the “voltage squared” effect for power reduction.

Since vSMP technology does not suffer from the penalties of cache synchronization and scheduling across cores running at asynchronous frequencies, it can deliver higher performance than architectures that use asynchronous clocking technologies.

Architectural Challenges and Solutions

vSMP architecture introduces several challenges, but several unique solutions have been created that address these challenges.

- **Switch time:** vSMP must ensure the process of switching between the Battery Saver CPU core and the main CPU core does not result in slower application load times and perceptible lags in user experience. To address this situation, NVIDIA implemented advanced circuits and logic to enable high speed switching efficiency. Internal simulations show that the total switching time, including the time to switch cores within

the chip and the time to stabilize the voltage rails (of the core that is turned on), is less than 2 milliseconds (ms) and this delay is not perceptible to end users.

- **Core Thrashing:** vSMP must prevent frequent back and forth switching between the Battery Saver and main cores when the workload varies around the threshold value used to trigger a switch between the cores, which would result in poor performance and negate any power savings benefits. To address this issue, sufficient intelligence and programmable hysteresis control is built into the CPU management algorithms that continually monitor and adapt to the workload, preventing any thrashing between the cores.

Power Benefits of vSMP

vSMP delivers significant power savings by minimizing both leakage power consumption in active standby state using the Battery Saver core, and dynamic power consumption at peak operating frequencies using the four main cores. Depending on the use case, vSMP dynamically enables and disables CPU cores to deliver the desired performance at the lowest possible power.

The chart below illustrates that Tegra 3 is lower power for all use cases. The chart measures Tegra 2 and Tegra 3 use, both on TSMC 40nm manufacturing technology.

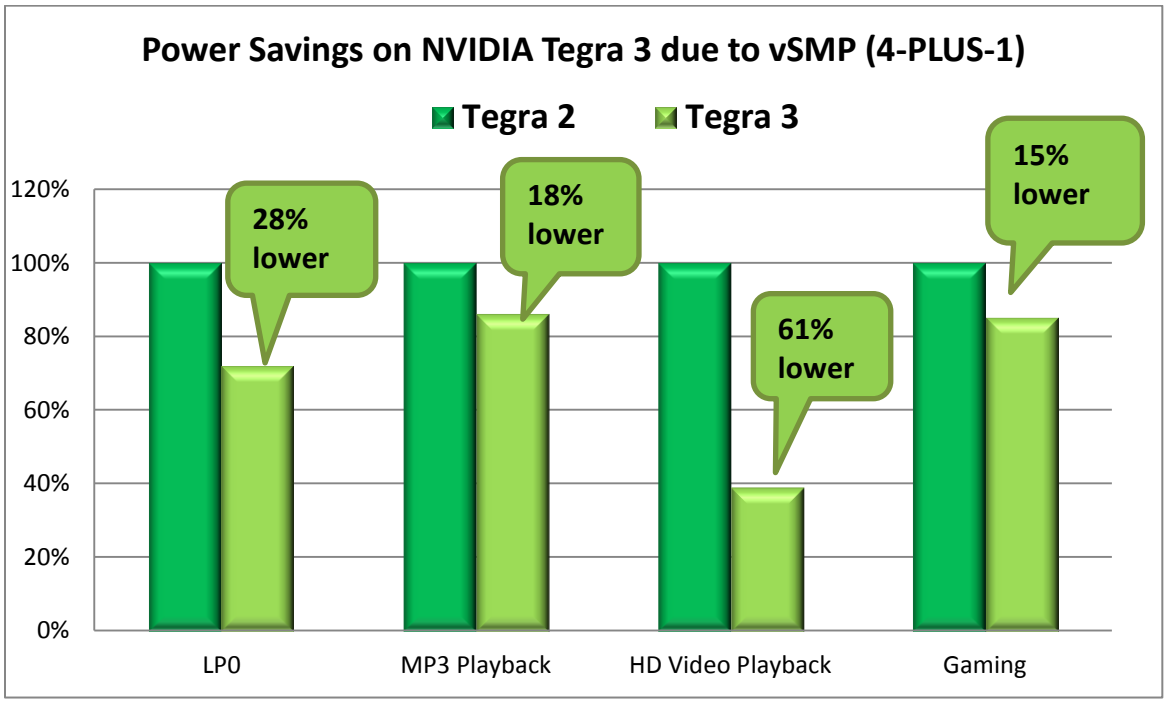

Figure 6 Power savings on NVIDIA Tegra 3 due to vSMP (4-PLUS-1)[1]

[1] Power measured as a sum of application processor power and DRAM power after normalizing for other system variables. LP0 is the lowest power state of the respective Tegra devices.

Power Benefits of Quad Core compared to Dual Core

In addition to vSMP, it's important to remember that more cores are better for power management than fewer cores. As an example, quad core CPUs deliver lower power at all performance points compared to dual core CPUs. The reason is that four cores can run at a lower frequency, and hence lower voltage when processing the same amount of work as a dual core CPU. Since power is proportional to the square of the voltage, the overall CPU power can be significantly less – and still complete the same amount of work.

Table 2 shows the measured power consumption and performance levels of NVIDIA Tegra 3 versus competing dual core processors running the Coremark benchmark, a popular mobile benchmark that measures single- or multi-core CPU performance. Note in the table below that Tegra 3's CPU consumes 2-3x less power than competing solutions when it is constrained to the same level of performance – i.e., each processor completing roughly 5k of Coremark "work". And even when Tegra 3 is run at a higher frequency, completing more than 2x the amount of Coremark "work", it is still less power than dual core solutions.

Mobile Processor	Measured Power (mW)[2]	Coremark Performance
Tegra 3 (each core running at 480 MHz)	579	5589[3]
OMAP4430 (each core running at 1 GHz)	1501	5673[4]
QC8660 (each core running at 1.2 GHz)	1453	5690[5]
Tegra 3 (each core running at 1 GHz)	1261	11667[6]

Table 2 Measured Power and Performance on Tegra 3 and Competing Processors

Note that Tegra 3 consumes lower power than competing dual core processors even when it has all four CPU cores running at 1 GHz. Because Tegra 3 uses fast process technology for its performance optimized CPU cores, these four cores are able to operate at higher frequencies using lower operating voltage ranges than competing processors. Since dynamic power consumption is proportional to the square of the operating voltage, Tegra 3achieves significant power savings even when operating at higher frequencies.

[2] **CPU power measured by taking total system power running Coremark (averaged over the full test) and subtracting system power at OS-Idle to arrive at CPU only power. Note that Tegra 3 runs in Battery Saver mode during OS-Idle. Measured on Tegra 3 reference design and competitive production devices.**

[3] CoreMark 1.0 : 5589 / GCC4.4.1 -O3 -mcpu=cortex-a8 -funroll-loops -falign-loops=8 -fgcse-sm -fno-tree-vectorize -marm / Heap / 4:PThreads

[4] CoreMark 1.0 : 5673 / GCC4.4.1 -O3 -mcpu=cortex-a8 -funroll-loops -falign-loops=8 -fgcse-sm -fno-tree-vectorize -marm / Heap / 4:PThreads

[5] CoreMark 1.0 : 5690 / GCC4.4.1 -O3 -mcpu=cortex-a8 -funroll-loops -falign-loops=8 -fgcse-sm -fno-tree-vectorize -marm / Heap / 4:PThreads

[6] CoreMark 1.0 : 11667 / GCC4.4.1 -O3 -mcpu=cortex-a8 -funroll-loops -falign-loops=8 -fgcse-sm -fno-tree-vectorize -marm / Heap / 4:PThreads

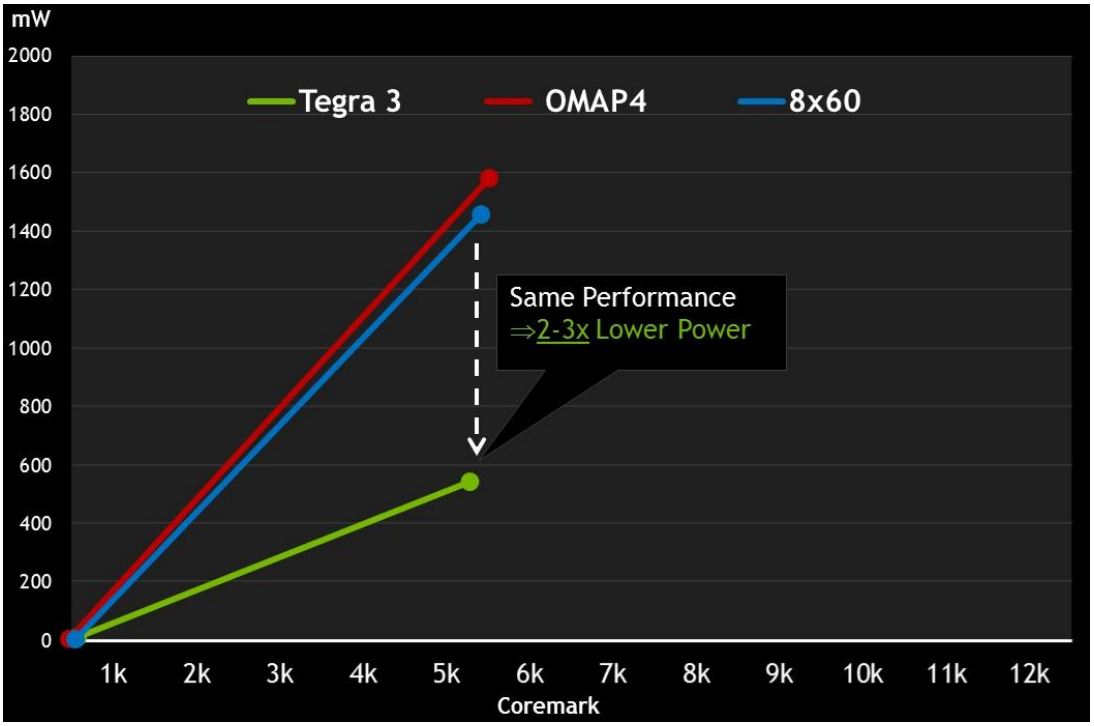

Figure 7 Tegra 3 power consumption when delivering the same performance as competing dual core processors[7]

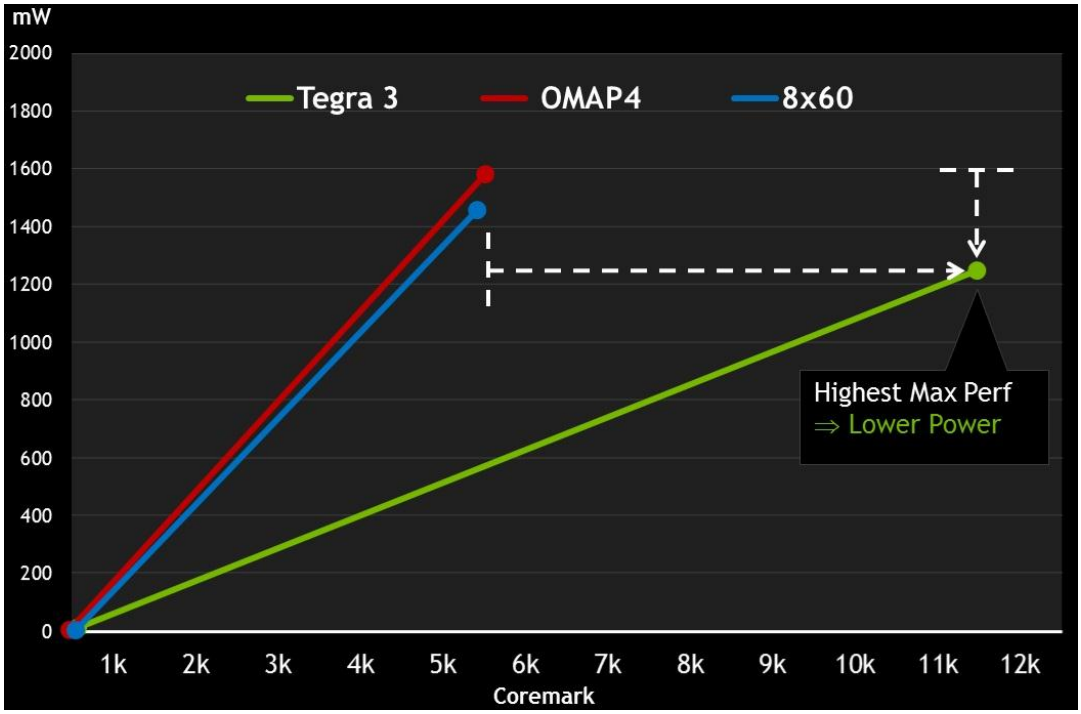

Figure 8 Tegra 3 power consumption delivering maximum quad core performance[8]

[7] Coremark scores reported in the graph were obtained using compile settings shown in Table 1 of the Appendix
[8] Coremark scores reported in the graph were obtained using compile settings shown in Table 1 of the Appendix

Conclusion

As the performance requirements of mobile applications increase, SoC vendors are not only adopting multi-core processor architectures to deliver the increased performance, they are also doing it to keep power consumption within mobile budgets. The Variable Symmetric Multiprocessing (4-PLUS-1) architecture employed in NVIDIA Tegra 3 delivers a new level of power savings that not only minimizes power consumption during active standby states, but also delivers quad core performance benefits while keeping dynamic power consumption within thermal budgets required for mobile devices. The use of the Battery Saver CPU core for background tasks and the use of main cores for performance-intensive tasks, enable NVIDIA Tegra 3 to deliver significantly lower power than competing mobile processors at all performance levels.

NVIDIA Tegra 3's 4-PLUS-1 quad core architecture will enable mobile devices to further push the performance envelope, and allow application and game developers to deliver new mobile experiences, all while extending battery life for the most popular use cases.

For information about the benefits of Quad core CPU in mobile devices please see the whitepaper titled "***The Benefits of Quad Core CPUs in Mobile Devices***".

Appendix

Coremark Compile Settings for Reported Scores
NVIDIA Tegra 3 (dual core mode, each core running at 1 GHz) CoreMark 1.0 : 5532 / GCC4.4.1 -O3 -mcpu=cortex-a8 -funroll-loops -falign-loops=8 -fgcse-sm -fno-tree-vectorize -marm / Heap / 4:PThreads
NVIDIA Tegra 3 (quad core mode, each core running at 1 GHz) CoreMark 1.0 : 11667 / GCC4.4.1 -O3 -mcpu=cortex-a8 -funroll-loops -falign-loops=8 -fgcse-sm -fno-tree-vectorize -marm / Heap / 4:PThreads
OMAP4430 (each core running at 1 GHz) CoreMark 1.0 : 5673 / GCC4.4.1 -O3 -mcpu=cortex-a8 -funroll-loops -falign-loops=8 -fgcse-sm -fno-tree-vectorize -marm / Heap / 4:PThreads
QC8660 (each core running at 1.2 GHz) CoreMark 1.0 : 5690 / GCC4.4.1 -O3 -mcpu=cortex-a8 -funroll-loops -falign-loops=8 -fgcse-sm -fno-tree-vectorize -marm / Heap / 4:PThreads

Table 3 Coremark Compile Settings for reported scores

Document Revision History

Revision Number	Notes
1.0	First Release
1.1	Updated Figure 6 on page 11 to correct Y-axis units and updated chart with revised power savings for Gaming.
1.2	Included Coremark compile configuration for scores reported in Table 2 and figures 7 and 8
1.3	Made updates to reflect 4-PLUS-1 naming

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, Tegra, 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