



Technical Brief

PowerMizer SX

Power Management Technology for
NVIDIA Notebook Chipsets

May 2006
TB-01873-001_v02

NVIDIA® GeForce® Go 6100/6150 and NVIDIA nForce® Go 430 notebook platform processors combine to provide the world's first notebook chipset with high-definition NVIDIA PureVideo™ technology, Microsoft® DirectX® 9 Shader Model 3.0 capability, high-definition audio processing, and enthusiast-class storage and communication features.

Most importantly, the NVIDIA Notebook Chipset Solution is designed from the ground up to deliver the longest battery life of any Turion-class chipset. This technical brief provides an overview of the chipset and system-level power management features known as NVIDIA PowerMizer® SX (System eXtensions) mobile technology.

PowerMizer SX Mobile Technology

NVIDIA has taken a new approach to power management with the PowerMizer SX system power management solution. PowerMizer SX keeps notebook power at a minimum, supplying power only when system activity demands it. System activity often occurs in quick bursts, with some units waiting for results or data from other parts of the system. For example, when transferring data from disk, the system is waiting for several microseconds to milliseconds while the data are being retrieved. These wait periods are opportunities to save power. NVIDIA PowerMizer SX keeps the system powered down during these wait periods, supplying energy only when data is ready to be transferred—in this example, from the disk.

PowerMizer SX controls the energy use with a dedicated System Management Unit (SMU) that continuously monitors system activity and events. This dedicated processor controls the power and speed of each functional unit in the chipset. Activity monitors in each unit record the activity and utilization levels. The SMU is alerted when activity levels change, permitting it to respond by adjusting the energy delivered to the unit. By modulating clock frequencies, data bus widths, and other parameters, the SMU can deliver the precise bandwidth that each device requires. In this manner, power is reduced while supporting the workload the device is being asked to handle. The real-time responsiveness of the SMU ensures that there is no compromise in performance when workloads increase.

Typical notebook power management schemes rely on a system processor response to interrupts to recognize and respond to power events. The finest event granularity is typically found in the milliseconds time scale (one thousandths). The dedicated SMU in NVIDIA nForce Go 430 can respond to power events with microsecond latency (1 millionth of a second). This allows power-saving mode changes to occur much more frequently. Many times a microsecond response can let system components reach lower power states, while millisecond response implies the chipset misses opportunities to reduce the power level before the next activity burst occurs. When reduced levels of activity occur, the system quickly returns to lower power-level states for greater power savings.

One example of power management is the ability to power down the drive controller while it is waiting for a hard disk to respond to a disk read request. Another example is the ability to reduce the power of the network controller while it is quiescent, and scale it up as soon as packets addressed to the notebook arrive.

Each functional unit has a preset performance based on the power management priority policy in effect. When operating on AC power, the SMU can adjust clocks to maximize performance while staying within the thermal design power (TDP) limits of the device. When the notebook is operating on battery power, each functional unit has its own device state that is optimized for longest battery life.

Programmable Thermal Management System

Maximizing performance within the thermal limits of the notebook is another innovation that NVIDIA is bringing to the market with the GeForce Go integrated graphics processor (IGP). This system enables the system designer to define the temperature that the chipset must not exceed, which ensures the IGP does not create a hot spot on the system, even when the chip is located close to the notebook skin. The GeForce Go IGP extracts the maximum possible performance without exceeding the maximum temperature allowed.

The PowerMizer SX thermal management system continually monitors the die temperature and combines this information with analysis of the system activity to stay just below the thermal design limit. This delicate balance is orchestrated by the combination of the SMU and the on-chip thermal sensing unit.

By implementing this in hardware using a programmable processor, GeForce Go IGP/NVIDIA nForce Go 430 ensures that thermal management is enforced even when the operating system hangs or becomes unresponsive.

Designed for Low Power

Active power management does help a chip run cooler. However, the following design considerations further reduce power.

- **Highly integrated design**

A highly integrated design allows maximum data to flow between system devices with minimum power dissipated. For example, with the IEEE 802.3 MAC integrated in the NVIDIA nForce Go 430, data crosses fewer system board level traces before reaching the host processor. Other chipset solutions with PCI Express–based network controllers use more power at both ends of the PCI Express link than is used by the short on-die communication path of NVIDIA nForce Go 430’s native network MAC.

Another example of the power-savings benefit of a highly integrated design is the clock synthesizer integrated in the NVIDIA nForce Go 430. External clock chips typically use more than 500 mW in addition to the power consumed by the Northbridge and Southbridge. When these PLLs are integrated into the GeForce Go IGP/NVIDIA nForce Go 430, the pair offers a flexible clocking solution that uses less power than discrete clock synthesizers. In addition, there is no inter-chip clock routing, which requires additional power-consuming I/O pads on each side.

- **Clock gating**

Clock gating extends the power savings within each functional block beyond the typical block-level clock-gating common to discrete system components. For example, an IEEE 802.3 MAC implemented in a discrete solution may only have its clock gated by an embedded system controller. If the device is on but idle, it will still consume power unnecessarily. The clock gating design within NVIDIA nForce Go 430 extends inside the controller so that, within the device, idle gates consume less power.

- **Low-voltage 90 nm process**

A low-voltage 90 nm process allows the GeForce Go IGP to run at industry-low voltage levels to get the cubic power savings offered by reduced voltage. NVIDIA has worked very closely with our foundry partner to optimize the 90 nm process for our mobile chips to achieve very low leakage. Chip designers have carefully analyzed the speed paths within GeForce Go IGP to maximize the use of low-leakage transistors on noncritical path circuits to further reduce power consumed.

- **Support for dynamic voltage**

Support for dynamic voltage allows the GeForce Go IGP to run at low voltage for longest battery life, or at higher voltage for ultimate system performance. Voltage transitions are managed in real time by the SMU for ultimate control.

System Power Management

Controlling the power used by the GeForce Go IGP/NVIDIA nForce Go 430 platform solution is only part of the overall effort to reduce power consumption and increase battery life.

- ❑ **Supports most efficient stutter mode**

Display refresh in a UMA-based AMD notebook requires the screen contents to be read from system memory 60 times per second. Even though the CPU may be totally idle, it is involved in reading screen data from system memory and transferring the data over the HyperTransport (HT) interface to the IGP. Stutter mode is a new chipset feature that allows the CPU core to reduce power from the C1 halt state (3.7 W) to the C3 state (1.7 W) and only keep the memory controller active for display refresh. In addition, the HT bus is allowed to disconnect, for lower power, in between display refreshes. GeForce Go IGP architecture provides a much higher percentage of time that the HT bus is disconnected, for a typical savings of 300 mW to 1.4 W compared to other Turion-class chipsets.

- ❑ **Supports FID/VID cycles for CPU P-state transitions**

- ❑ **Supports C0, C1, C2, C3, C1E, and C4 states**

- ❑ **Supports C3 Pop Up**

C3 Pop Up is a unique feature that allows coherent and/or noncoherent bus master activity managed by the chipset to occur without returning the CPU from C3 (1.4 W) to C0 (8.7 W). Historically, the ACPI model (based on the Intel FSB restrictions) has forced the CPU to return to C0 full-power operation when any bus master activity occurs in the system. Furthermore, to prevent accesses to the Intel CPU cache while the CPU is in the C3 state, the ACPI drivers have implemented methods to disable bus mastering. This has made it impossible for the system to save power when primarily performing CPU idle tasks, such as playing audio or transferring data to/from disk. The GeForce Go IGP/NVIDIA nForce Go 430 media and communications processor (MCP) platform has custom modes that permit devices to fetch DMA data from system memory without exiting the CPU from the C3 state. This allows the CPU to spend less time in C0, reducing overall system power consumption.

- ❑ **Supports AMD Turion Rev F Dual Core C1E state without any SMM code additions**

The GeForce Go IGP/NVIDIA nForce Go 430 MCP supports the new AMD Rev F C1E message and will automatically bring the CPU into C1E state without special software. The benefit of the C1E state is that the CPU only consumes C3 level power (1.4 W) instead of C1 level power (3.7 W) whenever it is in the halt state, resulting in significant power savings.

- ❑ **Superior suspend and hibernate power management**

When a notebook is put in suspend mode (ACPI S3- suspend to RAM), the system state is saved in system memory and the chipset powers down all other system devices to use as little power as possible. The memory is put into a sleep state where only a small amount of current is used to support self-refresh.

The NVIDIA nForce Go 430 MCP keeps a small core powered to respond to wake events such as a touch on the keyboard or movement of a mouse. During suspend, GeForce Go IGP is shut down completely and NVIDIA nForce Go 430 MCP draws an estimated 30 mW, ensuring a long-lasting battery in this mode. A notebook in suspend mode should be able to remain in that state for days. Several contemporary notebooks support less than a day's worth of suspend; this is unacceptable for users who want to leave their notebook in suspend over the weekend. When a wake event happens, the NVIDIA nForce Go 430 MCP responds immediately by supplying power and clocks to the rest of the system for a speedy resume.

In hibernate mode (ACPI S4-suspend to disk), the system state is saved to the hard drive and all system components are shut off except the MCP, which has to respond to wake events. In hibernate mode, NVIDIA nForce Go 430 MCP uses only 0.01 mW so that the notebook can stay in hibernate almost indefinitely. Once again, when the user opens the notebook lid or presses the power button, MCP springs to life, restoring power and clocks to the rest of the system for a fast resume experience.

- ❑ **Supports other standards**
Supports instantly available PC (IAPC), ACPI 2.0, PCI PM 1.1 standard power management APIs, and operating system managed power schemes.

I/O Power Management

HyperTransport

- ❑ Supports HyperTransport link disconnect and STOP/REQ protocol. HT Disconnect is optimized in GeForce Go IGP to allow the CPU's bus interface unit to be "disconnected" an estimated 72 percent of the time with a standard notebook display.
- ❑ HT Tri-State support in the GeForce Go IGP and NVIDIA nForce Go 430 allows the HT uplink and HT downlink to be tri-stated for reduced power.
- ❑ HT Low Frequency Operation saves power on both the HT uplink and the HT downlink.
- ❑ Adaptive HT Link resizing is supported by GeForce Go IGP and NVIDIA nForce Go 430 to adapt to use fewer lanes when power savings takes priority over performance.

PCI Express

- ❑ PCI Express is a very high-frequency system interconnect that uses more power than legacy system busses like PCI.
- ❑ Adaptive System Link Management (ASLM) of the ×16 graphics interface, supported by both GeForce Go IGP and NVIDIA notebook graphics processing units (GPUs), allows graphics to dynamically negotiate the number of PCI Express lanes to use for running graphics. The PCI express ×16 interface can be dynamically resized to operate with 16, 8, 4, 2, or 1 lane active to achieve the best balance of performance and power consumption.

- ❑ Dual ×1 PCI Express Peripheral Interface can be enabled or disabled independently to reduce power consumed.
- ❑ No battery life compromise USB 2.0 support; USB 2.0 controllers can be powered off individually. When a USB device is attached to most other chipsets, the CPU is forced from the C3 state (1.7 W) to the C1 state (3.7 W) by disabling the C3 mode. NVIDIA nForce Go 430 is unique in its support of simultaneous USB device attach and CPU C3 state.
- ❑ IEEE 802.3 Gigabit Ethernet MAC with Coma mode support is integrated in NVIDIA nForce Go 430 with fine-grained clock gating as well as block-level gating. Other chipsets rely on external Ethernet controllers that use the power-hungry PCI Express interface (which can consume 100 mW on each side of the connection), whereas the NVIDIA nForce Go 430 is directly connected to the HT bus via low-power on-die interconnect. Coma mode allows the Gigabit PHY to enter a low power state as if it were part of the clock gating tree of the MAC and to instantly spring to life when packets addressed to the MAC arrive.
- ❑ SATA 2 controllers with slumber mode support a high-performance 3 Gbps drive interface that can also operate at 1.5 Gbps for power savings. The controllers can be independently powered down. The NVIDIA nForce Go 430 SATA controller supports SATA drives with slumber mode that reduces SATA PHY power from 600 mW to 150 mW.

Graphics Power Management

True UMA Graphics

True UMA graphics does not require a local frame buffer for substantially lower system power operation than chipsets using a discrete local frame buffer memory. The GeForce Go IGP graphics core is optimized to provide the industry's best HT disconnect duty cycle for lower system power.

Without the need for a local frame buffer dedicated for graphics, system BOM cost and system power is reduced. The low demand on system memory for GeForce Go IGP's optimized display refresh uses less power than a dedicated frame buffer.

CineFX 3.0

The NVIDIA CineFX® 3.0 engine has dedicated GPU hardware that greatly offloads complex geometry calculations from the CPU. This helps CineFX 3.0 extend battery life (the CPU at full utilization is the biggest power consumer in a notebook PC), yet deliver superior performance-rendering cinematic-quality graphics. This ability directly translates to extended battery life for 3D gaming. By supporting both Vertex Shader 3.0 and Pixel shader 3.0 instructions in hardware, the GeForce Go IGP graphics processor is able to perform 3D graphics operations most efficiently. 3D engines using older shader technology need to use multiple passes to render complex shaders, thus use more power.

Clock for clock, the GeForce 6 graphics core has been proven more efficient than any other graphics architecture. This efficiency translates directly into extended battery life for demanding 3D applications.

PureVideo

The NVIDIA PureVideo™ processor is a revolutionary new programmable processor dedicated to delivering improved video playback quality and longer battery life. PureVideo technology offloads decompression of MPEG-2 and Windows Media Video from the CPU. The small PureVideo processor is architected specifically to run video algorithms and includes 50 new video instructions and uses fewer cycles for greater efficiency.

The PureVideo processor is a small, yet capable element of the graphics processor. Other chips use general-purpose shader hardware for video post-processing. Running video on the small PureVideo processor instead of the large 3D engine is far more efficient. PureVideo offloads video from the shaders, which translates to lower power and extended battery life for video playback.

SmartDimmer LCD Backlight Control

NVIDIA SmartDimmer™ LCD Backlight Control is a feature uniquely manageable through the IGP and is an intelligent way to manage notebook panel power consumption. The IGP driving the notebook panel is the only component fully aware of, and responsible for, driving the display on the panels.

This unique capability forms the basis for SmartDimmer technology, which lets customers preset brightness preferences through a Control Panel (Figure 1) and lets the IGP then manage the panel display within these limits. As a result, SmartDimmer can reduce power consumption in one of the most significant and always-on power expenders in a notebook—the display panel. In turn, this reduction seamlessly delivers enhanced battery life for the user.

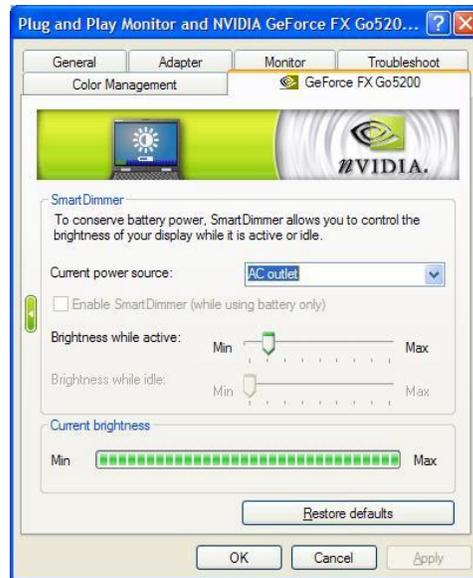


Figure 1. SmartDimmer User Controls

Summary

NVIDIA has delivered several generations of PowerMizer features in GeForce Go notebook GPUs. Now, PowerMizer SX mobile technology extends power management capabilities from the GPU to include chipset functions.

System power management includes a broad set of notebook subsystems such as system busses, storage, communications, CPU, LCD backlight, and more.



Notice

ALL 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 THE 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, CineFX, GeForce, NVIDIA nForce, PowerMizer, PureVideo, and SmartDimmer 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

© 2006 NVIDIA Corporation. All rights reserved.



NVIDIA.

NVIDIA Corporation
2701 San Tomas Expressway
Santa Clara, CA 95050
www.nvidia.com