Whitepaper

NVIDIA's Next Generation Notebook Technology:

Optimus™

V1.0

Table of Contents

Executive Summary

As the notebook industry continues to grow, consumers are trying to find the best compromise between battery life and performance. With the increasing popularity of HD media, gaming, and GPU powered applications, the number of GPUs included in notebooks continues to increase.

While integrated graphics has historically delivered good battery life, it has failed to deliver the graphics and GPU performance end users demand. As a result, a few years ago NVIDIA developed a technology dubbed “switchable graphics” which allowed the end-user to choose which display adaptor would be used. In essence, this technology brought the best of both worlds, as it offered the battery life of an integrated graphics solution and the performance of a discrete GPU. Unfortunately, there were limitations to the technology resulting in the end user having to execute a fairly involved procedure to harness the most from the platform, resulting in only 1% of users ever switching between the two graphics systems.

With NVIDIA’s new Optimus technology, users can now experience the full performance benefits of a discrete GPU with the battery life of an integrated graphics solution. NVIDIA Optimus automatically, instantaneously, and seamlessly optimizes the notebook to offer the best performance or best battery life depending on the application. And it simply works.

The Growth of Notebook PCs

Notebook PC growth has exploded in recent years. The desire to work and play anywhere has pushed the PC industry to deliver compelling new mobile products and technologies year after year.

Rapid advancements in notebook component design and capability have improved the mobile computing experience in many areas including increasing battery life, delivering higher performance, lowering power requirements, while providing smaller and thinner chassis. Notebooks can now include both local and wide-area wireless connectivity, and powerful notebooks are now available at lower price points.

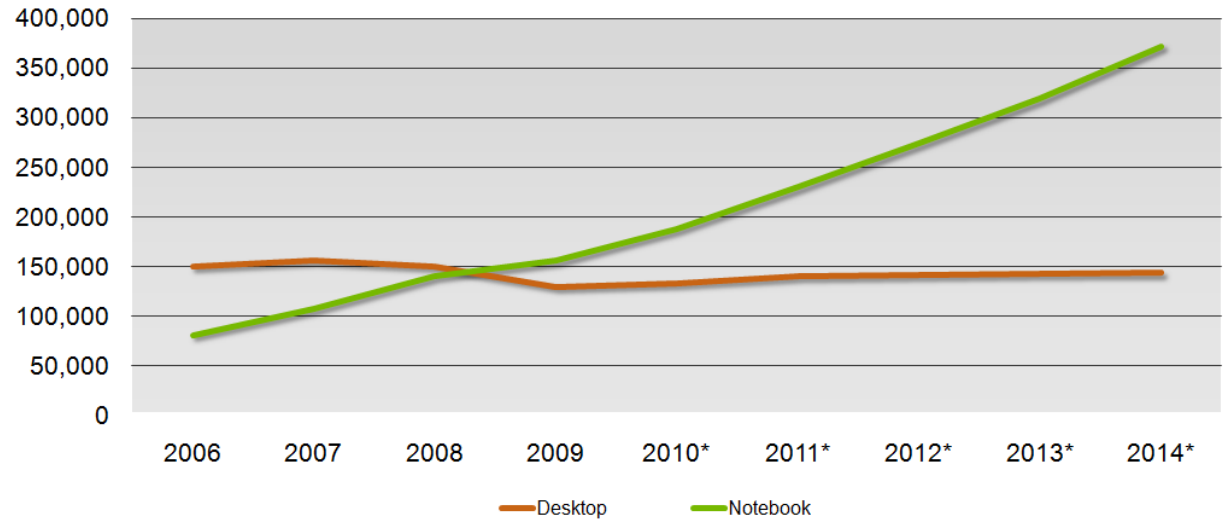

Figure 1: Notebook Sales Outpace Desktops

Notebook usage will continue to expand as mobile platforms continue to provide increasing performance and functionality, and users increasingly value untethered computing.

Importance of Battery Life

As notebook PCs have grown in popularity, the activities users perform with them have shifted. Notebooks were previously used as a mobile work desktop – Microsoft Excel, Word, PowerPoint, email, etc. Today “entertainment” is as important as “work” for many users, and more important for many others. The top activities now include multimedia consumption (watching movies, surfing the web, playing games), multimedia creation (sharing, editing, and storing photos and movies) and social media (Facebook, Twitter, MySpace and Flickr).

For years, notebook designers have strived to reduce size and weight, while increasing battery life and performance. This trend is still in full force as we are now seeing battery life in excess of 10hrs, desktop-level performance, and super thin and light notebooks designed of carbon fiber, anodized aluminum, zinc and magnesium composites, and other advanced materials. Battery life is of substantial importance. Retailers like Best Buy list battery life in their Sunday flyers.

System vendors promote battery life as a primary feature of their new notebooks. Intel markets "Ultra Low Voltage" CPUs that help increase notebook battery life and reduce notebook thickness ("z-height").

Industry benchmarks and publications have followed this trend. BAPCO's popular MobileMark benchmark, while not exactly representing expected battery life for all usage models, is a primary battery-life benchmarking tool for notebooks used at many publications today. Some publications have developed their own battery life tests comprised of scripts that automate Web surfing with and without wireless enabled, or play DVDs until the battery is depleted. These publications must demonstrate battery life in their reviews, and while their tests may not be as varied as MobileMark, or fully representative of typical notebook usage, they still provide useful comparative data. There is no doubt that end users and the industry care deeply about notebook battery life.

The Importance of Performance

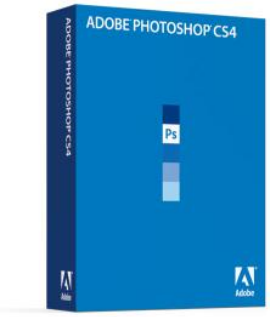

As mobile computing platforms become more commonplace, consumer expectations for their notebooks have changed. Users now expect a full desktop computing experience from their notebooks when watching movies, surfing the web, playing games, or updating social media Web pages. They have been trained for instant gratification, and watching an hourglass on their computer screen will not do. Dropped frames in video, poor response from games, and slow performance are not ideal.

But these types of applications require processing power, and as processing power increases, maintaining long battery life becomes much more difficult. Consumers want the untethered computing experience without the compromise to battery life that results from the increased demands placed on their mobile PCs.

The Mobile PC Dilemma: Performance vs Battery Life

The age old compromise for notebooks is performance vs battery life. Consumers are forced to prioritize performance or battery life, as one feature typically suffers significantly in order to accommodate the other.

Nowhere has the importance of mobile PC processing power been illustrated more clearly than with the rise of the netbook. A netbook was originally defined by its limitations:

> ***It's not a laptop; it's a netbook.***
>
> *"Based on a groundbreaking low-power microarchitecture, the Intel® Atom™ processor powers small, sleek Internet devices designed to go where you go. When all you need is to stay online so you can keep in touch with friends or follow your favorite Web sites, netbooks with the Intel Atom processor deliver convenience and flexibility in an incredibly small, amazingly smart package."*
>
> *Source:* ***Intel*** *(http://www.intel.com/Consumer/Learn/Internet-Devices/atom-netbook-detail.htm)*

Given the limitation of standard Atom-based netbooks, the result for consumers was disappointment and very high return rates, as much as 30% in some cases.

The introduction of NVIDIA's ION GPU redefined the entire netbook category and transformed expectations forever. With the addition of the NVIDIA ION GPU, the netbook could finally do what people wanted the platform to do. Whether the industry realized it or not, users were asking for a GPU. With the netbook platform's transformation as a profound proof point, the notebook industry is putting even more emphasis on including a discrete GPU on notebook motherboards.

Notebooks Need GPUs

Another notebook usage trend is the increasing desire to use multimedia capabilities. For instance, college students frequently use their notebooks as their primary TV and movie watching platform. This multimedia trend is leading to increased attachment of GPUs in mobile platforms, and it's happening worldwide, especially in developing regions.

A GPU not only provides the ability to play 3D games and watch crisp, vibrant, flawless high definition video, but also delivers new capabilities like GPU computing and GPU-accelerated Flash video playback. The top multimedia applications today leverage the GPU to deliver a better mobile experience to the end user.

While the increased performance and capability provides a more enjoyable notebook experience, it often conflicts with the other desired feature – better battery life. This isn't just a GPU phenomenon. Increased display brightness, faster CPUs, and higher RPM hard drives – anything that is high performance - all impact battery life.

To reduce mobile GPU power usage, NVIDIA has developed a number of patented technologies over the years that improve power efficiency while still maintaining GPU performance. Let's look at some of these technologies in the next section.

NVIDIA Battery Life Innovation

NVIDIA has been investing in notebook power efficiency technology for years and has made notable contributions to the industry.

NVIDIA Switchable Graphics

For many years the mobile PC industry has pursued the ideal mobile computer – one that has both great performance and long battery life. This often resulted in unique, and often crude, solutions that would improve either performance OR battery life, but did not deliver the combination of both great performance AND great battery life.

In one example, certain early notebooks permitted a user to choose the graphics system they desired to run (‘integrated graphics’ for low power or ‘discrete graphics’ for high performance), but had the unfortunate side effect of requiring the user to reboot the system for the switch to take place. A reboot was required because the operating system and graphics software (driver) were not aware of the switch, and the switch was managed by the system BIOS. For obvious reasons this solution was not widely adopted, but it did show the interest from PC builders to solve the problem.

NVIDIA launched a technology in 2007 which went a long way to solve this problem – Switchable Graphics– that dramatically changed the way a notebook performed and operated, by allowing the notebook to switch graphics subsystems without rebooting. Often this switch could take place in a matter of 5-10 seconds. For the first time a user could in theory have easy access to the best of both worlds – manually switch to integrated graphics for longer battery life, and manually switch to discrete graphics for higher performance – all in a single notebook, and all without rebooting. This was considered a groundbreaking milestone in mobile computing, and ushered in a new wave of notebook PCs. The technology was widely adopted by OEMs and received positive feedback from tech-savvy users and press.

Unlike the crude solutions of years prior, Switchable Graphics technology was primarily software driven. The graphics drivers and operating system enabled this capability. The rest of the PC graphics industry quickly followed and offered various versions of switchable graphics.

Unfortunately, while switchable graphics ushered in a new PC capability, it still had a few important usability and functionality problems which resulted in users rarely getting the desired benefit.

Switchable Graphics Technology

With the arrival of switchable graphics technology, users were able to reap the benefits of the discrete GPU when running intensive applications, and switch to the IGP (integrated graphics processor) during less intensive applications in order to preserve battery life. Functionally, the two separate graphics cores share the display using a hardware multiplexer that steers rendering and display duties between the graphics cores. The block diagram below depicts the logic flow of NVIDIA's switchable graphics technology.

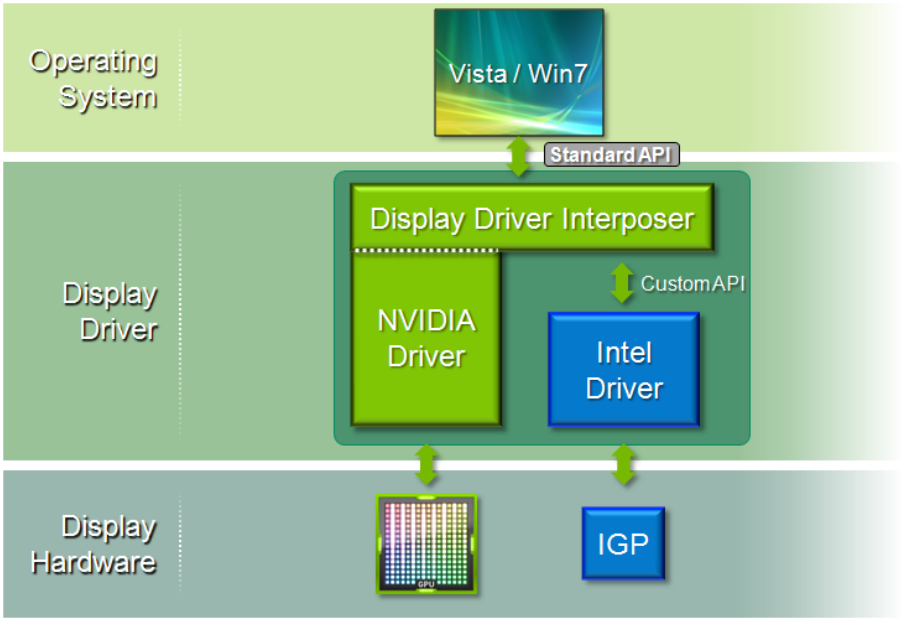

Figure 2: NVIDIA Switchable Graphics Data Flowchart

In order for the system to switch between the IGP and the discrete GPU, the user must manually change a setting within NVIDIA's Switchable Graphics Applet. This applet presents a simple menu that prompts the user to select a preferred graphics core based on a given workload. Once the selection is made, the system BIOS (SBIOS) exposes an interface that the Switchable Graphics driver uses to enable or disable the discrete GPU. In effect, the Switchable Graphics driver acts as a proxy display driver and shuttles OS information to the driver associated with the active and enabled display device.

Once the active graphics core (IGP or GPU) is ready to output a rendered scene or image, a hardware multiplexer (device which provides switching of multiple inputs) will switch to the correct display device as directed by the Switchable Graphics driver in order for the video output to be sent to the desired display.

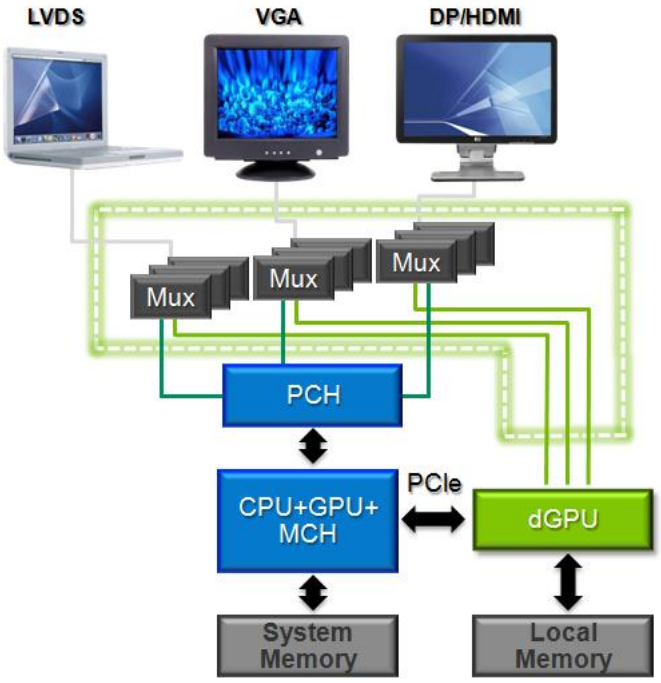

Figure 3: Switchable Graphics – Hardware Implementation

As you can see in Figure 3, the use of several multiplexers results in an extremely complicated layout. Multiplexers require additional space on the motherboard and dramatically increase trace lengths as each display interface needs to be routed twice. This results in a decrease in signal integrity and requires additional PCB layers which come with added cost and complexity.

The Issues with Switchable Graphics

Although NVIDIA's Switchable Graphics technology provides a dramatic increase in battery life for notebooks featuring discrete graphics, several hurdles and limitations are evident:

- **Manual Mode Changes** – Every time a switch between integrated and discrete graphics mode is desired, the user needs to manually change the setting. If users forget to enable the GPU before launching a game for example, they need to quit the game and enable the GPU before re-launching the application.

- **Transition time** – Due to the multiplexer and driver load and unload process, there is typically a delay of 5 to 10 seconds or more to switch between integrated graphics and the discrete GPU.

- **Blocking Applications** – Often applications "block" a mode switch. 3D and video applications and games like Solitaire and Minesweeper would cause a block. Each must be closed before the system will permit switching between the integrated graphics and the GPU. If users forget to close a 3D or video application, they often

would not realize the graphics mode did not change back to integrated graphics as expected.

- **Screen Flicker** – As the system switches between the integrated graphics and discrete GPU, the screen flickers due to the driver load time and multiplexers. This can be an annoyance or cause panic as users might assume there is a problem with the system.

- **Increased Cost** – Switchable Graphics relies on extra digital logic including multiplexers and additional display interfaces, therefore BOM cost is higher than a standard platform.

The result of all the above is that users rarely switched, because they rarely remembered the current graphics mode, preferred not to close their applications, and often didn't understand how to switch modes.

Another solution was needed that would allow the user to have the performance of the discrete GPU readily available when needed, while also being able to run on the integrated graphics for maximum battery life. Ideally, this would be done without user interaction and accomplished automatically in the background.

In order to overcome the above hurdles, NVIDIA developed a new technology which uses a combination of software and hardware innovations to automatically, seamlessly, and instantaneously deliver great performance and great battery life - **NVIDIA Optimus Technology.**

NVIDIA® Optimus Technology

NVIDIA Optimus Technology delivers great battery life and great performance, in a way that simply works. It *automatically* and *instantaneously* uses the best tool for the job – the high performance NVIDIA GPU for GPU-Compute applications, video, and 3D games; and low power integrated graphics for applications like Office, Web surfing, or email.

The result is long lasting battery life without sacrificing great graphics performance, delivering an experience that is fully automatic and behind the scenes.

When the GPU can provide an increase in performance, functionality, or quality over the IGP for an application, the NVIDIA driver will enable the GPU. When the user launches an application, the NVIDIA driver will recognize whether the application being run can benefit from using the GPU. If the application can benefit from running on the GPU, the GPU is powered up from an idle state and is given all rendering calls.

Using NVIDIA's Optimus technology, when the discrete GPU is handling all the rendering duties, the final image output to the display is still handled by the Intel integrated graphics processor (IGP). In effect, the IGP is only being used as a simple display controller, resulting in a seamless, flicker-free experience with no need to reboot.

When less critical or less demanding applications are run, the discrete GPU is powered off and the Intel IGP handles both rendering and display calls to conserve power and provide the highest possible battery life.

The beauty of Optimus is that it leverages standard industry protocols and APIs to work. From relying on standard Microsoft APIs when communicating with the Intel IGP driver, to utilizing the PCI-Express bus to transfer the GPU's output to the Intel IGP, there are no proprietary hoops to jump through. NVIDIA's new hardware and software design seamlessly blends into the existing frameworks.

How Optimus Differs from Switchable Graphics

Prior to the launch of Optimus, consumers looking for the highest performance and longest battery life obtained notebooks using switchable graphics technology. When an application needed higher performance, the user could turn on the GPU. When the user wanted to conserve battery power, the GPU could be disabled. When comparing the two technologies one can see substantial advantages provided by Optimus technology.

	Optimus	Switchable Graphics	Detail
Automatically Enable GPU for Intensive Applications			**Switchable graphics forces the user to manually change a setting.**
Automatically Disable GPU for Non-Intensive Applications			**Switchable graphics forces the user to manually change a setting.**
Near-Instant Transition Time			**Switchable graphics relies upon a multiplexer resulting in a delay when switching modes.**
No Screen Flicker During Transition			**Screen flicker occurs with switchable graphics when the new driver is being loaded.**
No Blocking Apps for Switching Modes			**3D games and applications need to be closed before switchable graphics will allow you to change modes.**
Supported Entirely by NVIDIA			**Switchable graphics requires cooperation and ongoing support from multiple vendors.**
Uses Standard Drivers			**Switchable graphics relies upon special drivers from each vendor.**
Doesn't Require Additional Notebook Components			**With switchable graphics, you need a multiplexer and additional display hardware adding cost and complexity.**
Uses Standard Industry APIs and Protocols			**Unlike switchable graphics which uses specialized drivers and extra components, Optimus relies on standard industry APIs and protocols like PCI-E and DirectX.**

Table 1: Optimus vs. Switchable Graphics Comparison

Optimus is Automatic

When using Switchable Graphics in the past, users were forced to manually enable the GPU before launching any application that needed additional GPU power. Users had to ensure no blocking applications were open before manually enabling the GPU. Once enabled, users could launch GPU-intensive applications. When the application was eventually closed, users would manually disable the GPU to conserve battery life. After the transition occurred, non-intensive applications could be launched.

In contrast, NVIDIA Optimus makes the process of enabling the GPU automatic. NVIDIA has done extensive testing and QA work to determine which applications should run on the GPU and has created profiles for these programs. Optimus can recognize CUDA GPU computing applications and many video playback applications that use DXVA GPU acceleration and will automatically switch to the GPU in those cases. As a result, users simply launch their desired application(s) and let Optimus take care of the rest. There are no additional steps or clicks required by the user in order to enable the GPU. The sample scenario illustrated below in Figure 4 depicts the launch process between the two solutions.

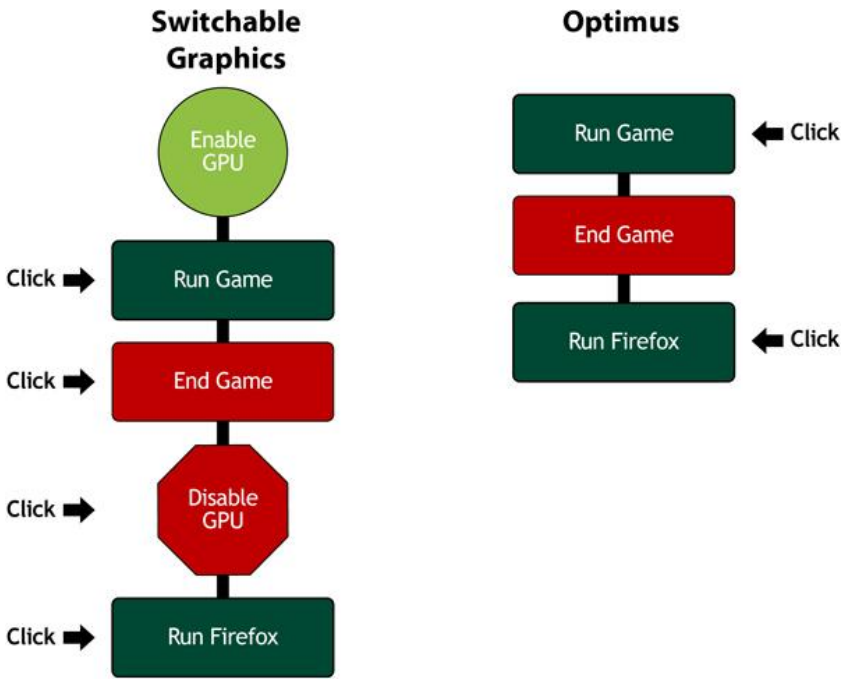

Figure 4: End-User Launch Process Comparison

An Overview of the Optimus Technology

As you will see in the detailed sections later in this whitepaper, NVIDIA's new Optimus technology has both software and hardware components. Developments for each area have allowed NVIDIA's engineers to make the transition between the GPU and IGP automatic. With Optimus, the user can stop worrying whether they should run on the IGP or GPU for a given application and simply focus on actually using the application.

When using non-taxing applications to accomplish basic tasks, like checking email or creating a document, Optimus recognizes that the workload does not require the power of the GPU. As a result, Optimus completely shuts off the GPU (and associated PCIe lanes) to provide the highest possible efficiency and battery life. In this case, illustrated in Figure5, the IGP will be used for all processing duties and will also act as a display controller to output the frames to the display.

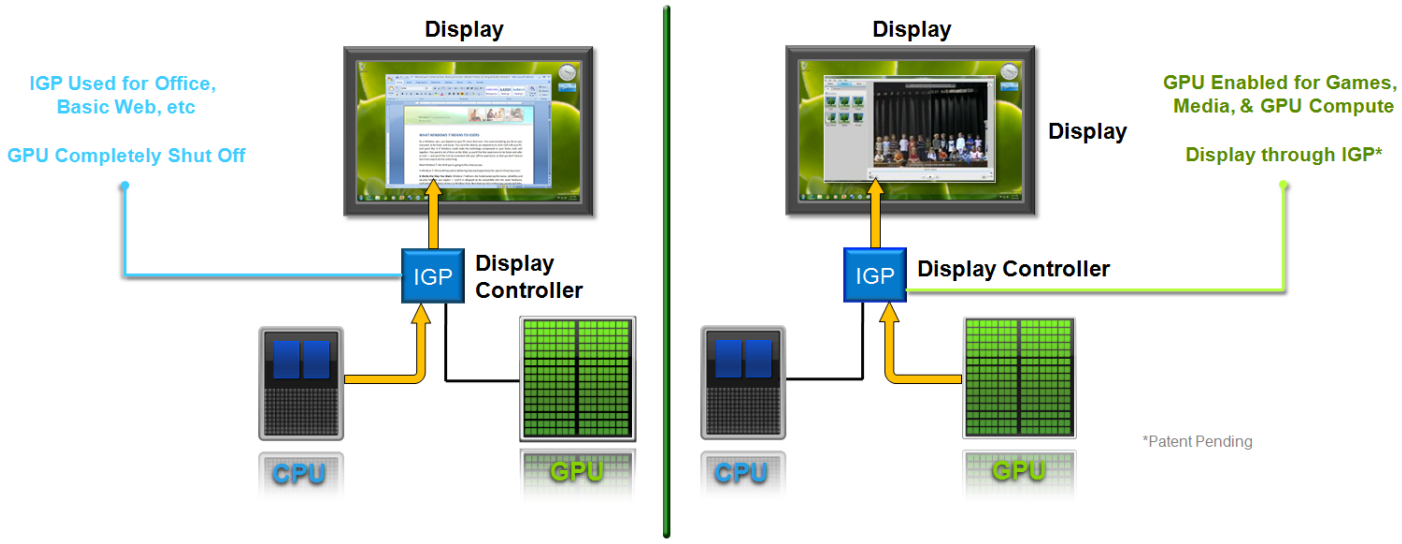

Figure 5: Application Scenario Comparison

As soon as applications that can benefit from the power of the GPU are invoked, like watching Flash video, gaming, or converting video from one format to another using CUDA, Optimus instantly enables the GPU. As shown in the figure above, the GPU handles all processing duties and the IGP is only used as a display controller to render the GPU's output to the display.

Optimus can be broken down into two key categories:

- **Optimus Software**
  - Optimus Routing
  - Optimus Profiles
- **Optimus Hardware**
  - System Advantages
  - Optimus Copy Engine

Optimus Software

Although several new hardware features are designed into NVIDIA GeForce 200M, 300M and Fermi GPUs which make Optimus possible, NVIDIA's software team has also implemented innovative technology to enable Optimus to function seamlessly without having to manually change settings.

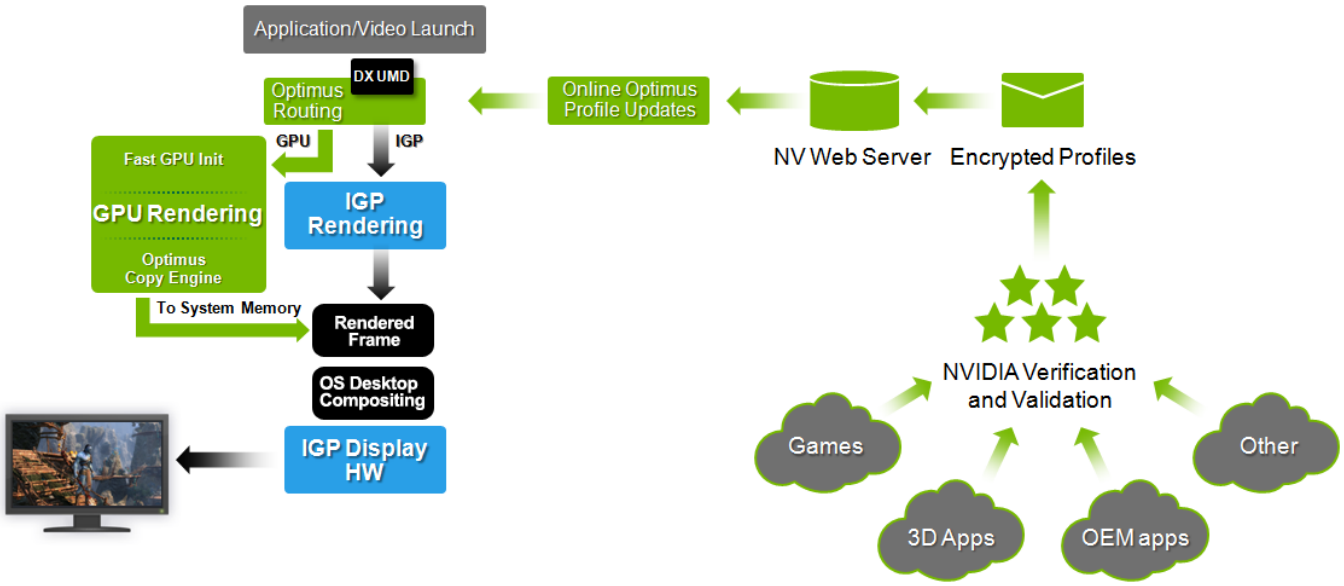

Figure 6: NVIDIA Optimus Flow

Optimus Routing Layer

Within the hardware interface layer of the NVIDIA GPU driver, the Optimus Routing Layer provides intelligent graphics management. The Optimus Routing Layer also includes a kernel-level library for recognizing and managing specific classes and objects associated with different graphics devices. This NVIDIA innovation performs state and context management, allocating architectural resources as needed for each driver client (application). Note that each application is not aware of other concurrent application usages of the GPU in this context-management scheme.

The allocation and management of architectural objects (hardware resources) provides several capabilities for compatibility management. The Optimus Routing Layer establishes a channel between the hardware and a driver client in a manner that minimizes the driver overhead for ascertaining hardware resources.

By recognizing designated classes, the Optimus Routing Layer can help determine when the GPU can be utilized to improve rendering performance. Specifically, it sends a signal to power-on the GPU when it finds the following three call types:

- **DX Calls**: Any 3D game engine or DirectX application will trigger these calls
- **DXVA Calls**: Video playback will trigger these calls (DXVA = DirectX Video Acceleration)
- **CUDA Calls**: CUDA applications will trigger these calls

Optimus Profiles

Another keystone of NVIDIA's Optimus technology is the ability to use a list of GPU-powered applications, each with its own unique Optimus Profile. NVIDIA has profiled applications based on whether the GPU can add quality, performance, lower power, or functionality. In doing so, Optimus is able to avoid powering up the GPU unless it is necessary.

Before an application has an Optimus Profile, NVIDIA runs the application through a series of verification and validation processes to ensure the highest quality and experience. Once validated, the updated Profiles are hosted on an NVIDIA web server and then automatically pushed out to the end user. Should the user not want automatic updates, they can disable this feature within the NVIDIA Control Panel.

As new applications and games are introduced in the market they may not have an Optimus Profile preinstalled, so NVIDIA has created a unique solution-- an NVIDIA-hosted Web server will push updated Optimus Profiles to the end user. Users also have the ability to customize Profiles and create new ones.

Similar to the SLI profiles found in NVIDIA drivers over the last several years, a section within the NVIDIA Control Panel displays the application-specific settings for Optimus. Although there is an extensive list of existing application profiles with the preferred graphics processor for each application, users are also able to add new application profiles and customize existing profiles from the list. The drop menu and associated settings for doing so are illustrated below in Figure 7.

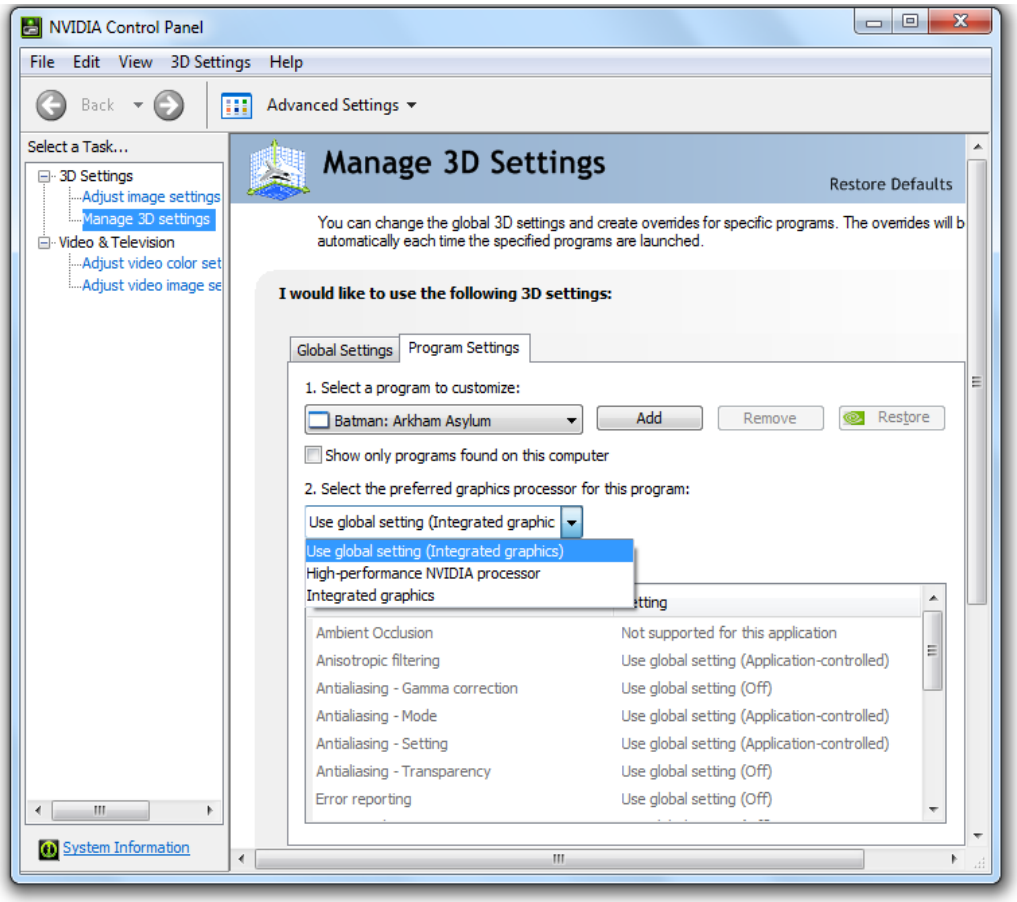

Figure 7: Application Profile Settings in NVIDIA Control Panel

Optimus Leverages Windows 7

Optimus technology leverages Windows 7's ability to allow two independent graphics drivers to be active at the same time. The standard Intel graphics driver is used along with the NVIDIA driver because both display adapters operate independently. Looking within the Windows Device Manager, you'll see two display adapters listed even if Optimus has turned the GPU off.

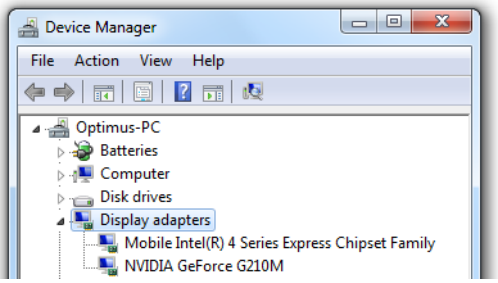

Figure 8: Device Manager Showing Two Active Display Adapters

Optimus Hardware

System Advantages

One of the primary advantages of Optimus compared to Switchable Graphics is that it does not need additional hardware like multiplexers. An Optimus implementation's simplicity reduces costs, eases design and verification effort, and speeds time to market. As noted earlier, all display connections are made directly to the IGP. Regardless of which display adapter is processing the workload, the IGP will push the output to the display.

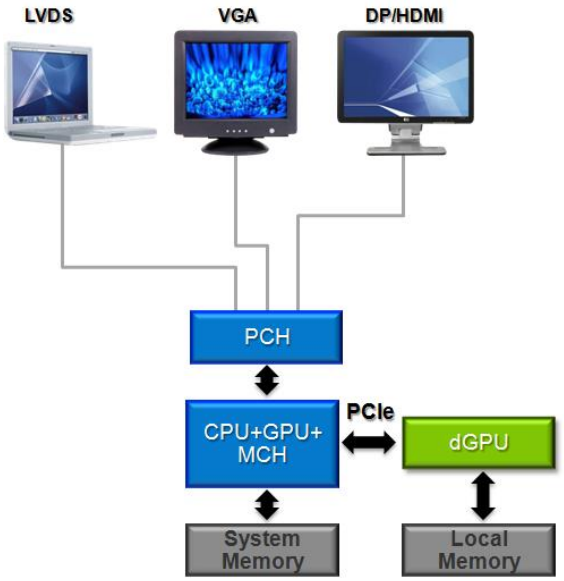

Figure 9: Optimus – Hardware Implementation

Optimus Copy Engine

In order to utilize the software innovations of Optimus, the GPU requires a creative new hardware feature, known as the Optimus Copy Engine.

Optimus avoids usage of a hardware multiplexer and prevents glitches associated with changing the display driver from IGP to GPU by transferring the display surface from the GPU frame buffer over the PCI Express bus to the main memory-based framebuffer used by the IGP. This occurs when an Optimus Profile indicates the application requires the GPU to be enabled. The key to performing the display transfer without negatively impacting 3D performance is the Optimus Copy Engine.

The Optimus Copy Engine is a new alternative to traditional DMA (Direct Memory Access) transfers between the GPU framebuffer memory and main memory used by the IGP.

Traditionally, mem2mem DMA transfers are performed by the 3D engine. To preserve coherency, the 3D engine is blocked from rendering until the mem2mem transfer completes. This time-consuming (synchronous) DMA operation can stall the 3D engine and have a negative impact upon performance. The new Optimus Copy Engine relies on the bidirectional bandwidth of the PCI Express bus to allow simultaneous 3D rendering and copying of display data from the GPU framebuffer to the main memory area used as the IGP framebuffer. (The IGP then reads the display data out of the framebuffer and sends it through a display interface (DVI, HDMI, etc) to the display). This asynchronous DMA operation provides data coherency and allows for a dramatic increase in both performance and efficiency.

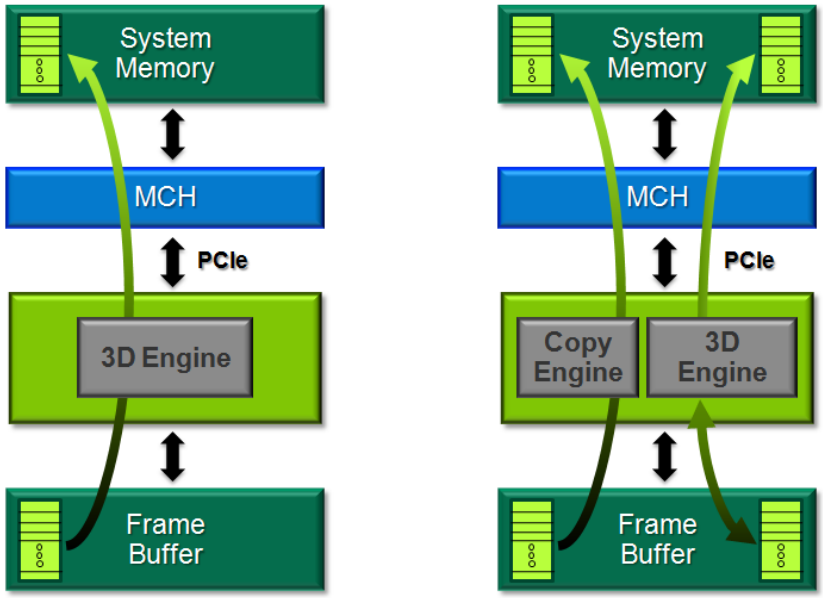

Figure 10: Readback Comparison

Sample Scenarios

Given the fact that Optimus technology is completely seamless, users can rest assured they are experiencing the highest possible performance without any sacrifice of battery life. Some sample scenarios are illustrated below to highlight the benefit the discrete GPU can have upon performance as well as depict when Optimus deems the integrated graphics as being sufficient.

Scenario 1 – Playing World of Warcraft

When gaming, you want to be able to load a game and have it run the way the developers intended. With an NVIDIA-based notebook, you're leveraging the power of the GeForce GPU with the stability of NVIDIA Verde notebook drivers and quality of NVIDIA's The Way It's Meant to be Played program.

Experience on IGP: Given the lack of processing power, framerates will be extremely low resulting in choppy gameplay and difficulty keeping pace with others online. In-game image quality settings need to be at very low settings in order to run, meaning you're missing out on the game's intended look and feel. In addition, the use of a non-NVIDIA driver means there's potential for a variety of compatibility issues and stability problems which can ruin the gaming experience.

Experience on GPU: When running on the GeForce GPU, the game will run at exceptional framerates and offer the end user fluid playback. Given the surplus of graphics horsepower, in-game settings and resolution can be raised to provide much better visuals and showcase the true potential of the game. Since you're running NVIDIA Verde notebook drivers, you're guaranteed to have the best compatibility, stability, and performance.

How Optimus Reacts: Optimus recognizes that a game has launched and appropriately enables the GPU for the highest performance. Should the game be paused for an extended period of time, the GPU will be shutdown to conserve power until the game resumes playing.

Scenario 2 – Convert and Copy Video to Zune HD

Systems which do not feature a discrete GPU are forced to rely upon the CPU for tackling heavy compute applications like video transcoding. If you're trying to convert and copy a video onto your portable device (ie: Zune HD) you'll have a much better experience when using a discrete GPU.

Experience on IGP: When no discrete GPU is present, you're forced to use a program which relies upon the CPU for processing power. For example, transcoding a video will take a considerable amount of time when running on CPU versus a GeForce GPU.

Experience on GPU: Systems with a GeForce GPU can take advantage of CUDA applications like Badaboom to harness the parallel processing power of the GPU for transcoding. Badaboom will transcode a video in a fraction of the time it takes to process on the CPU.

How Optimus Reacts: Optimus detects that a GPU Compute application has launched and powers up the GPU for the highest performance. As soon as the CUDA application is closed, the GPU will be turned off to conserve power.

Scenario 3 – Web Browsing

When simply browsing the Web and checking email, the processing power of the GPU is largely unutilized. However, as Flash-based content continues to grow there are situations like watching streaming HD content where Optimus can harness the GPU's power to provide the best possible experience.

Experience on IGP: Running on integrated graphics, general Web browsing will perform as expected. This will provide the highest possible battery life without sacrificing the browsing experience.

Experience on GPU: Although the GPU will provide the same experience as the integrated graphics when doing general Web browsing, the GPU will consume more power resulting in lower overall battery life.

How Optimus Reacts: Optimus detects that there are no intensive applications launching and keeps the GPU powered down. In doing so, the system is able to observe the highest possible battery life without any sacrifice of performance or the overall user experience. However, if the user is streaming Flash video using a website like YouTube or Hulu, Optimus will recognize the performance and quality benefits the GPU provides to Adobe Flash 10.1 (especially HD and high quality content) and will enable the discrete GPU.

Conclusion

Notebooks are developed in a wide variety of sizes and shapes in an effort to provide the ideal balance of battery life and performance for varied consumer needs. As a result, consumers are forced to prioritize performance or battery life, as one feature typically suffers significantly in order to accommodate the other.

With NVIDIA Optimus, things simply just work. Users can enjoy watching a HD movie, surfing the Web or playing top 3D games with the assurance of getting the best performance and the longest battery life.

Document Revision History

V1.0

- Initial release

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
© 2009 NVIDIA Corporation. All rights reserved.