

White Paper
David Galus
Graphics Platform
Application Engineer
Kirk Blum
Graphics Solution
Architect
Intel Corporation

Hybrid Multi- monitor Support

Enabling new usage
models for Intel[®]
Embedded Platforms

January, 2010



Executive Summary

Hybrid Multi-monitor is one form of Intel's Hybrid Graphics Support where integrated graphics (graphics built into the motherboard chipset) is available to operate simultaneously with add-in or "external" graphics. In the past this use case was not supported by Intel, but now Hybrid Multi-monitor is supported with Intel's Graphics Drivers, starting with the latest Intel chipsets incorporating integrated graphics (e.g., Intel® 4 and 5 Series Platforms). Please refer to your chipset capability list, or contact your Intel Representative to verify your chipset has this capability. This feature enables concurrent function of Intel's integrated Graphics Processing Unit (GPU) along with a discrete GPU solution, allowing for operability of greater than two independently-driven displays. With the availability of these additional display options, innovative embedded usage models may be employed, while creating the possibility of a more immersive and efficient system design. This paper will focus on describing embedded usage models supported by this feature, modes of display operation, Hybrid Multi-monitor enablement, operating system (O/S) considerations, functionality and limitations, and provide a power analysis for this scenario.

Benefits:

- More simultaneous independent displays.
- "Full value" from the Intel platform solution by enabling simultaneous use of integrated graphics with external graphics.
- Provides greater graphics scalability and flexibility.
- Enables a wider range of graphics based usage models.




The Intel® Embedded Design Center provides qualified developers with web-based access to technical resources. Access Intel Confidential design materials, step-by step guidance, application reference solutions, training, Intel's tool loaner program, and connect with an e-help desk and the embedded community. Design Fast. Design Smart. Get started today. www.intel.com/embedded/edc.



Contents

Introduction	5
Business Challenge	5
Purpose of Document	5
Hybrid Multi-monitor Definition	5
Embedded Usage Models Enabled.....	5
Intel® Integrated Graphics Basic Display Architecture	7
Intel® Graphics Modes of Operation	10
Dual Display Clone	10
Dual Display Twin	10
Extended Desktop/Dual Independent Head.....	10
Hybrid Multi-monitor Enablement	10
Requirements	11
Enablement Process – Windows*	11
Functionality	13
Enablement Process – Linux*	14
General Limitations	14
Power Analysis.....	15
Conclusions	16



Introduction

Business Challenge

As technology advances and environments become more interactive and immersive, the need arises for additional electronic display capabilities. Most systems that provide integrated graphics capabilities are limited to two simultaneous independent displays. Using discrete “add-in” or “external” GPU’s to add more displays in the past meant the integrated graphics was disabled thus losing the ability to use that capability. Hybrid Multi-monitor functionality addresses this limitation.

Purpose of Document

The purpose of this document is to detail Intel’s solution for Hybrid Multi-monitor functionality describing: embedded usage models supported by this feature, basic display architecture, modes of display operation, Hybrid Multi-monitor enablement, O/S considerations, functionality and limitations, and power analysis for this scenario.

Hybrid Multi-monitor Definition

Hybrid Multi-monitor (also known as ‘Multi-GPU Multi-monitor’) enables concurrent function of Intel’s integrated GPU along with a discrete GPU solution, allowing for operability of greater than two independently-driven displays. Intel’s Integrated GPU alone provides support for two independent displays. With the addition of a discrete GPU (depending on the manufacturer and model) additional independent displays may be added to the system, creating the ability to drive multiple (more than two) independent displays. Details on particular operational modes and benefits for embedded customers of using these multiple displays are explained in later sections.

Embedded Usage Models Enabled

Having the ability to drive greater than two independent displays opens up new opportunities for embedded customers. Scenarios, for instance, which required the use of four displays for the same piece of equipment, can now all be operated on a single system. One segment where this may be applicable is the Gaming Segment. Casino gaming machines typically use multiple displays, and having the ability to expand past the typical two display capability allows for less system complexity and more efficient system management.



Another segment of potential interest is Digital Signage. A modern signage concept may employ many displays. Imagine, for example, a flight arrival board in an airport terminal. These setups typically span eight to ten displays. Therefore, having the ability to control more than two displays from one computer system should make this setup easier to manage.

See [Figure 1](#) and [Figure 2](#) for graphics detailing the previously mentioned setups. These examples are just a sampling of what Intel believes may be of interest to our embedded customers; however, the configuration possibilities are endless.

Figure 1. Gaming Environment



Source: <http://www.daylife.com/photo/07Hz2x4eyqd1c>



Figure 2. Digital Signage Environment



Intel[®] Integrated Graphics Basic Display Architecture

Hybrid Multi-monitor is a mode of operation where both an integrated GPU and discrete GPU are used concurrently to drive two or more active displays across GPUs, thus giving the user an extended view ([Figure 3](#)). Each GPU uses both its own render and display engines. Each GPU creates, renders and scans out from its own display buffers, corresponding to a unique frame buffer with unique settings for resolution, refresh rate, color depth, scaling, and rotation ([Figure 4](#)).

The O/S will handle control of the multiple GPU display adapters appropriately ([Figure 5](#)). For example, Windows* XP supports The Microsoft* Windows XP Display Driver Model (XPDM) which allows loading and support of multiple graphics drivers. Windows* 7 continues that legacy XPDM support but also adds WDDM v1.1 which, like XPDM, allows for simultaneous multiple graphics drivers (Windows* Vista WDDM v1.0 did not allow this capability). Operating system applications will be adapter-unaware through use of the O/S GUI APIs



and will utilize the adapter associated with the primary display, regardless of which display the image is located on.

Note: Some applications may be adapter-aware, e.g., full-screen applications and system applications like the compositor. A number of software tools designed to assist multi-monitor use are available from third parties. One example is the UltraMon* utility for multi-monitor systems, which helps with the position of applications, assists desktop wallpapers and screen savers in multi-monitor configurations.

Hybrid Multi-monitor mode is recommended to be accomplished using a discrete third-party PCI Express* graphics card either into the PEG slot of the platform or into an available PCI Express* slot routed off of the I/O subsection of the chipset. See the implementation section for additional possible configurations.

The BIOS will enumerate the integrated GPU and discrete GPU graphics adapters and no changes should be required in the third-party vendor graphics drivers assuming they were originally designed to coexist with multiple graphics drivers.

Figure 3. Extended Desktop Model

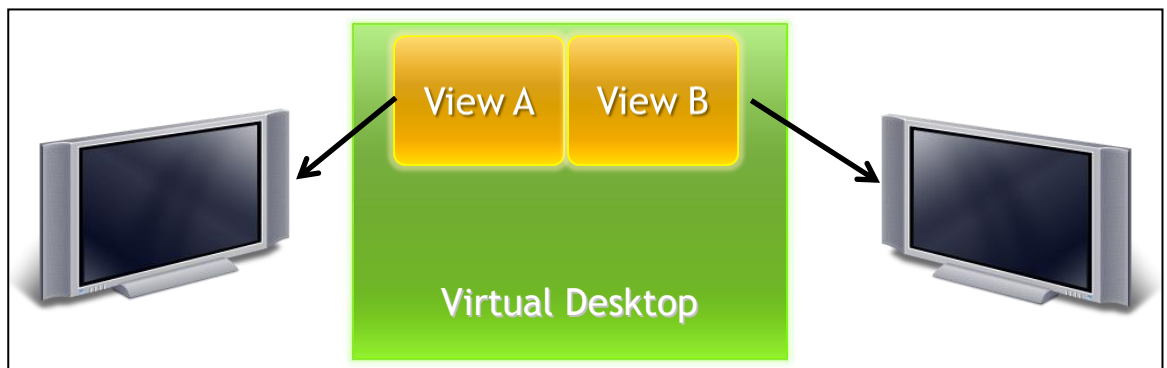




Figure 4. Hybrid Multi-monitor Model

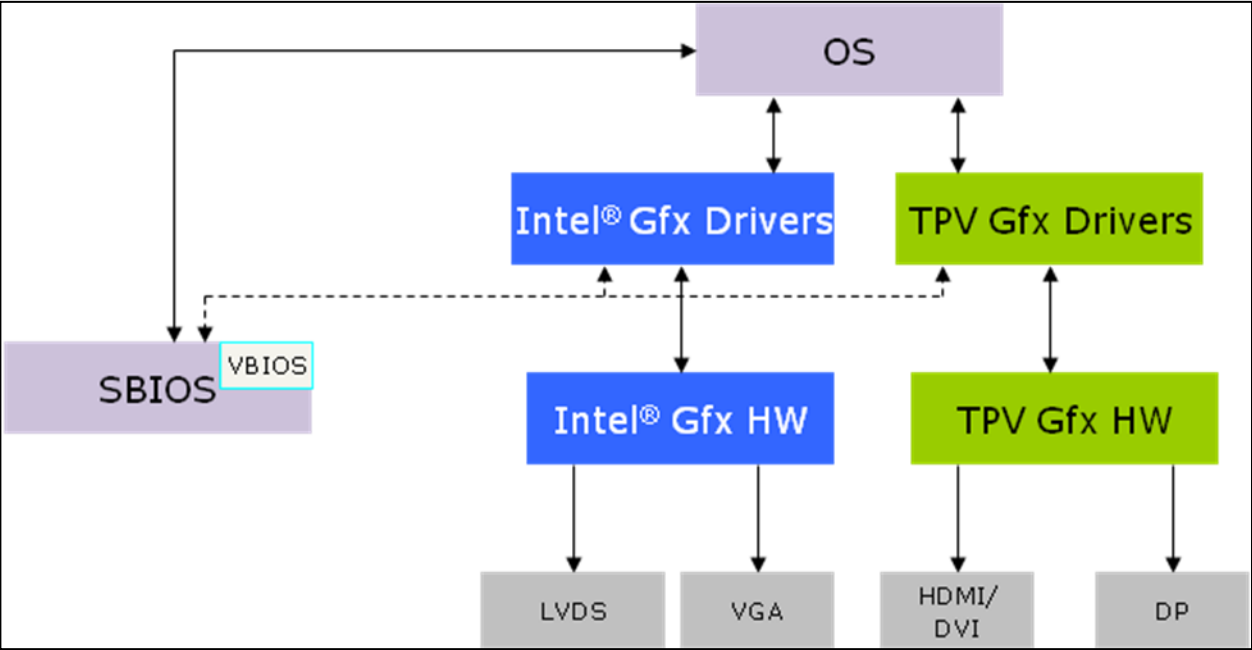
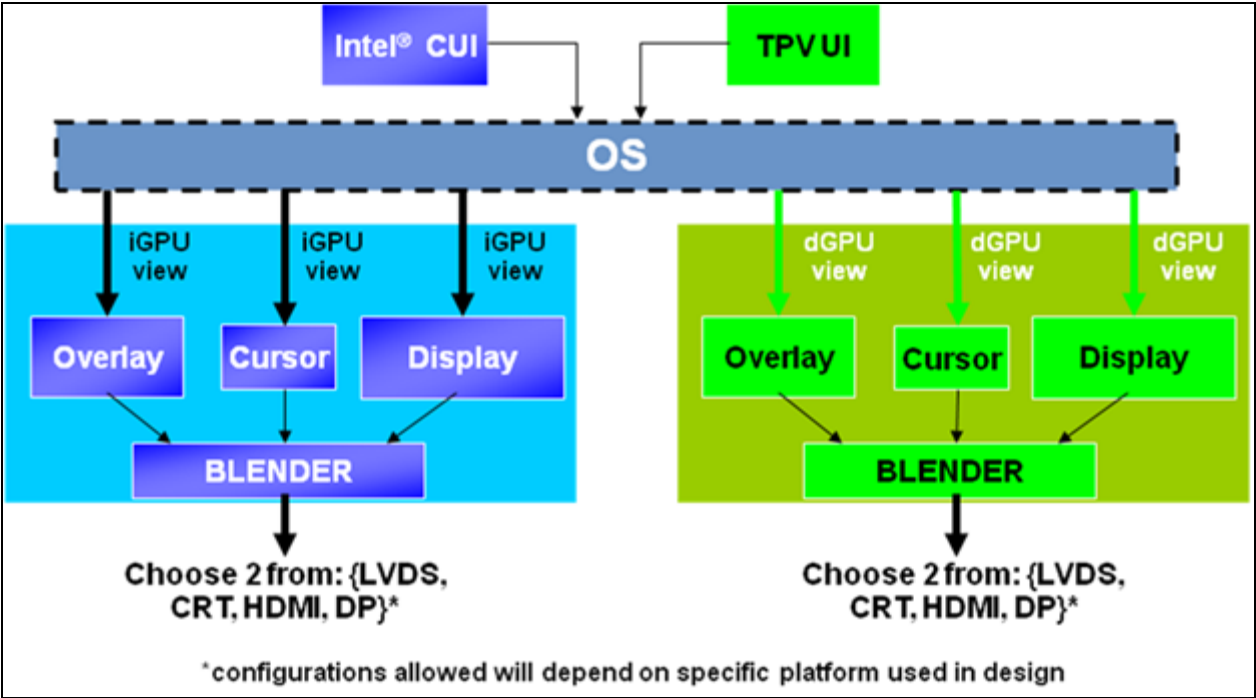


Figure 5. Hybrid Multi-monitor O/S Interaction





Intel[®] Graphics Modes of Operation

The graphic modes described below will be supported on displays controlled by Intel's integrated GPU.

Dual Display Clone

This mode is also known as Multi-pipe Simultaneous Output and refers to a graphics-operating mode where two displays showing the same image are driven by two output pipes. The resolution and refresh rate are independently selectable on each display. The primary display's content and color depth are duplicated on a secondary device.

Dual Display Twin

This mode is also known as Synchronous Output and refers to a graphics-operating mode where two displays showing the same image are driven by a single output pipe. The primary display's content, resolution, color depth, and refresh rate, are duplicated on a secondary device. This mode is being phased out of Intel's Scalable and Low Power product lines.

Extended Desktop/Dual Independent Head

DIH mode enables two display pipes to drive two independent display planes simultaneously on two separate displays ([Figure 3](#)). When in DIH, each independent display mode list is created and available for each display. Please note that available display modes may be directly impacted depending on platform configuration and available bandwidth. Windows implementation of DIH is called Extended Mode where the DIH displays are logically connected to extend the desktop across the multiple displays.

Hybrid Multi-monitor Enablement

Enablement of Hybrid Multi-monitor is not very complex; however, there is a specific procedure to this process that must be followed. Unfortunately, there also may be particular platform limitations that may affect your ability to enable this capability, e.g., BIOS restrictions. This section describes the Requirements, Enablement Process, Functionality, Limitations and Power Analysis for Hybrid Multi-monitor support.



Requirements

Because of Linux*/Windows XP/Windows 7 native support for heterogeneous display adapters, the additional requirements of the system are minimal and primarily are there to ensure Hybrid Multi-monitor will not become disabled.

Besides running using one of Intel's Hybrid Multi-monitor supported platforms, (e.g., Intel® 4 and 5 series products) containing integrated graphics, the system in consideration will need to have the ability to support a discrete GPU, either in an available PEG, PCIe* or PCI slot. The system BIOS should allow the user to select between setting the integrated GPU or discrete GPU as the primary boot display. For Hybrid Multi-monitor support please set the integrated GPU as the primary boot display.

Enablement Process – Windows*

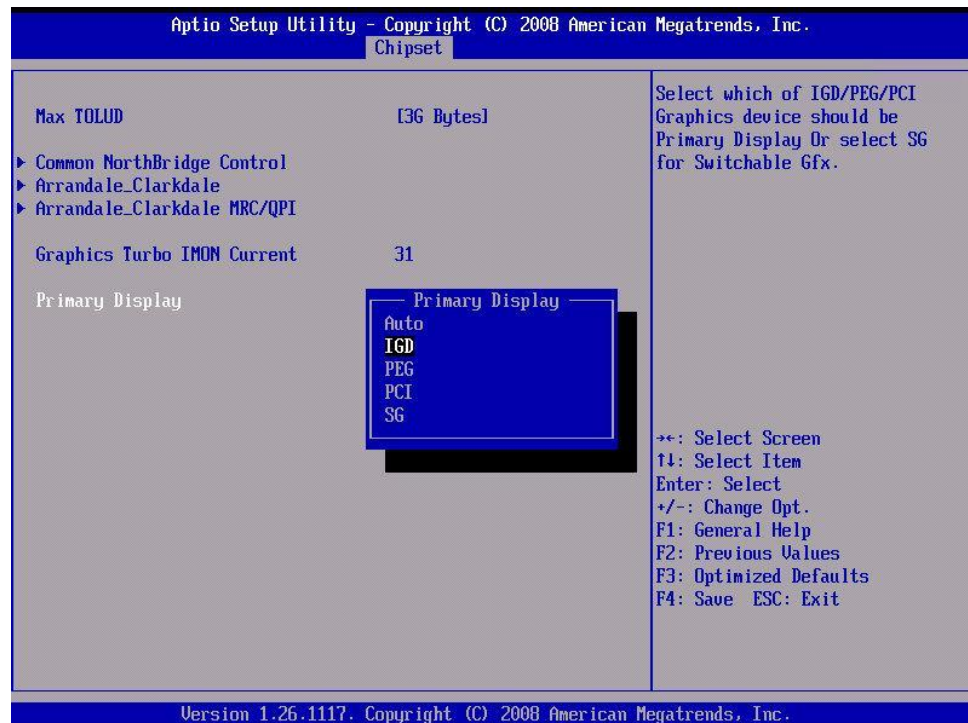
1. Begin this process without the discrete GPU seated in the system.
2. Enter the system BIOS and select **IGD** as the Primary Graphics Device. This selects the "Integrated Graphics Display/Device".

Note: The screenshot in [Figure 6](#) is from an Intel® 5 Series platform. Typically this feature will be found under the Chipset features for the "Northbridge" even though with modern solutions, the Northbridge only exists in concept.

Note: System BIOS vendors must not disable the secondary GPU, so that the O/S can manage support of this device.



Figure 6. BIOS Enablement



3. Boot into the O/S (Windows XP / Windows 7).
4. Install drivers requested by O/S for the Intel® integrated GPU (refer to platformsw.intel.com or downloadcenter.intel.com for the latest Intel graphics driver release).
5. Shut down the system, insert the discrete GPU into an available PEG, PCIe or PCI slot and connect displays accordingly.
6. Boot into O/S.

Note: It may be necessary to boot to Windows* **Safe** mode or **VGA Enabled** mode and install vendor drivers for each device separately. However, if this is not done, just ensure that the Intel® Graphics Device is set as primary within the system BIOS setup.

7. Install the drivers as requested by O/S for the discrete GPU.

Note: Some third-party graphics drivers have additional steps to enable all display outputs available on their card. Please refer to the instructions or third-party help desk for details on this.

8. Set the Windows Display properties as referenced below.




Functionality

Adding a discrete GPU solution to a system, in conjunction with an integrated GPU allows for Hybrid Multi-monitor support. This will enable support for greater than the two displays supported by Intel integrated graphics on supported platforms.

As mentioned in a previous section, O/S applications will be adapter-unaware, and the GPU used during system operation will be determined by the setting of 'Primary Graphics Device.'

Once enabled, Windows Display Properties **Settings** can be used to manage the heterogeneous display adapters and configure the multiple displays appropriately. See [Figure 7](#) below.

Figure 7. Windows* Display Setup





Enablement Process – Linux*

1. Enable in the system BIOS as mentioned above in the Windows Enablement Process.
2. Configure xorg.conf for the integrated graphics.
3. Verify the system boots and can display X Desktop.
4. Install external graphics card and configure per the instructions for the external card.
5. Use Xinerama* to enable extended mode across all of the displays if you do not want the displays running independently.

General Limitations

When implementing Hybrid Multi-monitor into your system design there are limitations to consider. These limitations are:

- Because Hybrid Multi-monitor requires the use of an external GPU, keep in mind the space constraints of your system, as greater space will be required with the discrete card.
- Power management and PCIe States versus graphics card power management – for enabling power management features of the PCIe bus, choose a discrete card that supports power management and ASPM.
- PCI/PCIe slot width available on board – Intel typically validates using x16 PCIe external graphics cards in the PEG slot. Depending on the embedded system design and power trade-offs, x1 x2 x4 or x8 widths may be possible. See the associated Platform Design Guide for details.
- PCIe lane width – depending on the system BIOS provided, this PCIe lane width may or may not be scalable; Intel typically validates using a x16 width.
- If a platform is designed with an Intel® chipset with integrated graphics capability and implements an on board external graphics solution, the integrated graphics might not be able to be used depending on the design.
- Discrete GPU form factor and MXM revision – the embedded products encompass many differing form factors, including motherboard down devices. If a particular form factor is used, ensure that ACPI functionality is supported through BIOS.
- BIOS enumeration and set up options – as a discrete display adapter will be used, verify that the system BIOS will enumerate this additional device.
- Primary display device selection (integrated/discrete GPU) – users will have a choice to either set the integrated/discrete graphics as Primary. Intel requires setting the integrated GPU as Primary, or, for mobile products, whichever device uses LVDS is preferred to be set as Primary.



- User Interface allows for selection of active monitors and adjusts settings – this is an optional API that may be developed, however Intel recommends using the User Interface delivered by the Operating System to manage these settings.
- Hot-Plug support – where supported, Hybrid Multi-monitor has not been validated with Hot-Plug support; however, expectations are that the primary display adapter should be maintained. This may be achieved either using native O/S support or by using a separate user-provided application.
- Behavior in S3/S4 power transitions – S3 and S4 transitions are the most problematic and may be affected by the quality of the external graphics driver(s). Although validated alone, Intel® graphics drivers may not have been fully validated with Hybrid Multi-monitor and S3/S4 power states. In particular, there is impact to S3 resume time.
- Mode enumeration and persistence along with PCIe bandwidth for performance considerations.
- Error handling and fail-safe mode when one of the devices does not function properly will be greatly affected by the third party driver and the operating system. It is recommended that the user reboots the system in this scenario.
- Microsoft* DirectX allows for 3D and media applications of only the primary to be enumerated to applications – when the user launches any 3D or media application, the O/S will assign the Primary GPU to this application. The user may need to avoid doing “drag and drop” of these applications to displays driven by the other GPU’s, as it may lead to unexpected behaviors, i.e., crashes, blank screens, display corruption, etc.
- TV Out – Certain Intel® integrated graphics products may not be able to drive a digital interface (HDMI/DVI/DP/sDVO) along with any integrated TV-out while a PCIe graphics card is connected. Please refer to specific product documentation to determine system limitations.
- Clone and Twin are only possible within a single GPU. It is not possible to Clone or Twin between integrated and discrete graphics.
- Rotation – It is recommended to set all displays to non-rotated before adding or removing other displays. Transitioning into and out of S3/S4 may cause rotation persistence issues.

Power Analysis

Power measurements were completed to understand the difference between the two operating modes, i.e., integrated displays alone versus use of both integrated and discrete displays. Based on the power data collected, while using integrated displays along with discrete displays, CPU power increased by around 1W when compared to using integrated displays alone. The tests were conducted on a single platform and the number should be used as guidance only.



The platform configuration and test setup is shown below:

- Intel® Core™ i7 Processor Fab C Board w/ (Intel® Core i7 SV 2133 MHz C0 Stepping, Intel® QM57 Express Chipset B3 Stepping)
- System BIOS Rev. 0.77 Video BIOS Rev. 1891
- 2 DDR3 1066MHz w/ECC 1GB SDRAM DIMMs
- Windows XP SP3
- Chipset Driver 9.1.1.1013
- GMA 14.41.2.5179 Driver
- NVIDIA* Display Driver v195.62
- EVGA NVIDIA GeForce 9800 GT 512MB DDR3 External Graphics Card
- 1 VGA (Integrated), 1 DP (Integrated), 2 DVI (Discrete)

The graphics were stressed using 3DMark06 developed by Futuremark* Corporation.

Conclusions

Hybrid Multi-monitor enables new visually immersive scenarios requiring the need for multiple monitors without the need for disabling integrated graphics. This allows for additional system cost savings in the final product. The Hybrid Multi-monitor feature may be enabled through specific steps followed with the latest Intel® chipsets incorporating integrated graphics, e.g., Intel® 4 and 5 Series Platforms. For additional design assistance or technical resources related to this topic please refer to <http://intel.com/embedded/edc>.

The Intel® Embedded Design Center provides qualified developers with web-based access to technical resources. Access Intel Confidential design materials, step-by-step guidance, application reference solutions, training, Intel's tool loaner program, and connect with an e-help desk and the embedded community. Design Fast. Design Smart. Get started today. <http://intel.com/embedded/edc>.



Authors

David Galus is a Graphics Platform Application Engineer with Intel's Embedded Communications Group

Kirk Blum is a Graphics Solution Architect with Intel's Embedded Communications Group

Acronyms

API – Application Peripheral Interface

BIOS – Basic Input/Output System

DDR – Double-Data-Rate

DIH – Dual Independent Head

DP – DisplayPort

DVI – Digital Visual Interface

ECC – Error Correction Codes

GPU – Graphics Processing Unit

GUI – Graphical User Interface

HDMI – High Definition Multimedia Interface

IGD – Intel Graphics Driver

O/S – Operating System

PCI – Peripheral Component Interconnect

PCIe – PCI Express

PEG – PCI Express Graphics

SDRAM – Synchronous Dynamic Random Access Memory

sDVO – Serial Digital Video Output

SV – Standard Voltage

VGA – Video Graphics Array

WDDM – Microsoft* Windows* Display Driver Model

XPDM – Microsoft Windows* XP Driver Model



INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL® PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER, AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT. Intel products are not intended for use in medical, life saving, or life sustaining applications.

Intel may make changes to specifications and product descriptions at any time, without notice.

This paper is for informational purposes only. THIS DOCUMENT IS PROVIDED "AS IS" WITH NO WARRANTIES WHATSOEVER, INCLUDING ANY WARRANTY OF MERCHANTABILITY, NONINFRINGEMENT, FITNESS FOR ANY PARTICULAR PURPOSE, OR ANY WARRANTY OTHERWISE ARISING OUT OF ANY PROPOSAL, SPECIFICATION OR SAMPLE. Intel disclaims all liability, including liability for infringement of any proprietary rights, relating to use of information in this specification. No license, express or implied, by estoppel or otherwise, to any intellectual property rights is granted herein.

BunnyPeople, Celeron, Celeron Inside, Centrino, Centrino logo, Core Inside, Dialogic, FlashFile, i960, InstantIP, Intel, Intel logo, Intel386, Intel486, Intel740, IntelDX2, IntelDX4, IntelSX2, Intel Core, Intel Inside, Intel Inside logo, Intel. Leap ahead., Intel. Leap ahead. logo, Intel NetBurst, Intel NetMerge, Intel NetStructure, Intel SingleDriver, Intel SpeedStep, Intel EP80579 Integrated Processor, Intel StrataFlash, Intel Viiv, Intel vPro, Intel XScale, IPLink, Itanium, Itanium Inside, MCS, MMX, Oplus, OverDrive, PDCharm, Pentium, Pentium Inside, skool, Sound Mark, The Journey Inside, VTune, Xeon, and Xeon Inside are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the U.S. and other countries.

*Other names and brands may be claimed as the property of others.

Copyright © 2010 Intel Corporation. All rights reserved.