

DECLARATION OF NATHANIEL E FRANK-WHITE

1. I am a Records Request Processor at the Internet Archive. I make this declaration of my own personal knowledge.
2. The Internet Archive is a website that provides access to a digital library of Internet sites and other cultural artifacts in digital form. Like a paper library, we provide free access to researchers, historians, scholars, and the general public. The Internet Archive has partnered with and receives support from various institutions, including the Library of Congress.
3. The Internet Archive has created a service known as the Wayback Machine. The Wayback Machine makes it possible to browse more than 450 billion pages stored in the Internet Archive's web archive. Visitors to the Wayback Machine can search archives by URL (i.e., a website address). If archived records for a URL are available, the visitor will be presented with a display of available dates. The visitor may select one of those dates, and begin browsing an archived version of the Web. Links on archived files in the Wayback Machine point to other archived files (whether HTML pages or other file types), if any are found for the URL indicated by a given link. For instance, the Wayback Machine is designed such that when a visitor clicks on a hyperlink on an archived page that points to another URL, the visitor will be served the archived file found for the hyperlink's URL with the closest available date to the initial file containing the hyperlink.
4. The archived data made viewable and browseable by the Wayback Machine is obtained by use of web archiving software that automatically stores copies of files available via the Internet, each file preserved as it existed at a particular point in time.
5. The Internet Archive assigns a URL on its site to the archived files in the format `http://web.archive.org/web/[Year in yyyy][Month in mm][Day in dd][Time code in hh:mm:ss]/[Archived URL]` aka an "extended URL". Thus, the extended URL `http://web.archive.org/web/19970126045828/http://www.archive.org/` would be the URL for the record of the Internet Archive home page HTML file (`http://www.archive.org/`) archived on January 26, 1997 at 4:58 a.m. and 28 seconds (1997/01/26 at 04:58:28). The date indicated by an extended URL applies to a preserved instance of a file for a given URL, but not necessarily to any other files linked therein. Thus, in the case of a page constituted by a primary HTML file and other separate files (e.g., files with images, audio, multimedia, design elements, or other embedded content) linked within that primary HTML file, the primary HTML file and the other files will each have their own respective extended URLs and may not have been archived on the same dates.
6. Attached hereto as Exhibit A are true and accurate copies of screenshots of the Internet Archive's records of the archived files for the URLs and the dates specified in the attached coversheet of each printout.

7. Attached hereto as Exhibit B are true and accurate copies of the Internet Archive's records of the archived files for the URLs and the dates specified in the attached coversheet of each file.
8. I declare under penalty of perjury that the foregoing is true and correct.

DATE: December 10, 2022

Nathaniel E Frank-White
Nathaniel E Frank-White

EXHIBIT A

<https://web.archive.org/web/20120426234300/http://www.ece.ucdavis.edu/vcl/pubs/2011.12.techreport.techscale/>

Toward More Accurate Scaling Estimates of CMOS Circuits from 180 nm to 22 nm

[Aaron Stillmaker](#)

[Zhibin Xiao](#)

[Bevan M. Baas](#)

VLSI Computation Laboratory

Department of Electrical and Computer Engineering

University of California, Davis

Abstract:

With deep submicron technology nodes other methods are needed to obtain scaling factors rather than the traditional scaling factors which held for the pre-submicron era. This work presents scaling factors between major technology nodes between 180 nm and 22 nm operating at voltages from 1.8 V to 0.7 V. Common operating data for these technologies were taken from the International Technology Roadmap for Semiconductors (ITRS). HSpice simulations that rely on the Predictive Technology Model (PTM) for transistor characteristics were used to find the scaling factors.

Paper



[PDF \(494 KB\)](#), (c) Copyright, 2011, University of California, Davis

Reference

Aaron Stillmaker, Zhibin Xiao and Bevan M. Baas,
"Toward More Accurate Scaling Estimates of CMOS Circuits from 180 nm to 22 nm,"
Technical Report ECE-VCL-2011-4 VLSI Computation Lab, ECE Department, University of California, Davis, December, 2011.

BibTeX Entry

```
@techreport{astill:techscaling,  
  author = {Aaron Stillmaker and Zhibin Xiao and Bevan Baas},  
  title = {Toward More Accurate Scaling Estimates of CMOS Circuits from 180~nm to 22~nm},  
  institution = {VLSI Computation Lab, ECE Department, University of California, Davis},  
  year = 2011,  
  month = dec,  
  number = {ECE-VCL-2011-4},  
  note = {\url{http://www.ece.ucdavis.edu/cerl/techreports/2011-4/}}  
}
```

[VCL Lab](#) | [ECE Dept.](#) | [UC Davis](#)

Last update: Dec. 11, 2011

<https://web.archive.org/web/20020210130046/http://www.vsi.org/members.htm>



go to

join a DWG

members only

news

events

library

related sites

glossary

about VSIA

join VSIA

member list

more info

home



Member List

[0-In Design Automation](#)
[3 DSP](#)
[A Priori Microsystems](#)
[Adaptive Silicon](#)
[Advantest](#)
[Alatek](#)
[Alcatel](#)
[Alcatel Internetworking](#)
[Altera](#)
[Amphion](#)
[Analog Devices](#)
[Antrim Design Systems](#)
[Arc Cores](#)
[Arasan Chip Systems](#)

[ARM](#)
[Artec Design Group](#)
[Artest Corp.](#)
[Artisan Components Inc.](#)
[The Athena Group](#)
[Atmos](#)
[Axeon Ltd.](#)
[Axyx Design Automation](#)
[Beach Solutions](#)
[BOPS, Inc.](#)
[Cadence Design Systems](#)
[Canadian Microelectronics](#)
[Cannon](#)
[Cast](#)
[ChipIdea Microelectronics](#)
[Chip Implementation Center](#)

[Circuit Semantics](#)
[Clearspeed Technology](#)
[Co-Design Automation](#)
[Conexant Systems](#)
[ControlNet](#)
[CoreEI](#)
[CG-CoreEL Programmable Solutions Ltd](#)
[CoWare](#)
[Cypress Semiconductor](#)
[Denso Corp](#)
[Design and Reuse](#)
[DNP Corp USA](#)
[Dolphin Integration](#)
[Easics NV](#)

[ECSI](#)
[Ellipsis Digital Systems](#)
[Elixent Ltd.](#)
[Embedded Solutions](#)
[ETRI Microelectronics Technology Laboratory](#)
[Fluence Technology](#)
[Forschungszentrum Informatik an der Universitat Karlsruhe](#)
[Fraunhofer-IMS Institute](#)
[Frontier Design](#)
[Fujitsu Ltd. Global UniChip](#)
[Goya Technology](#)
[HCL Technologies India Pvt](#)
[HDLab, Inc.](#)
[Hewlett-Packard](#)
[Hitachi Ltd](#)
[IBM Microelectronics](#)

[IDEC \(IC Design Education Center\)](#)
[IMEC](#)
[IMMS](#)
[Improv Systems](#)
[Industrial Technology Research Institute](#)
[Infineon Technologies](#)
[InSilicon](#)
[Institute for System Level Integration](#)
[Intel](#)
[Intellitech Corp](#)
[InTime Software](#)
[IPTC Corporation](#)
[Jennic Ltd.](#)
[Jet Propulsion Laboratory](#)
[Korea Electronics Technology Institute](#)

[Lateral Sands](#)
[Logic Research](#)
[LogicVision](#)
[LSI Logic](#)
[LSI Systems Inc](#)
[Macronix International](#)
[Malardalen University](#)
[Matsushita](#)
[MCHIT](#)
[Memec Core](#)
[Mentor Graphics](#)
[Microelectronics - Harbin Institute of Technology](#)
[Microelectronics Research Institute PROGRESS](#)
[MIPS Technologies](#)
[Mitsubishi Electric](#)
[Mixel](#)
[Modelware, Inc.](#)
[Montery Design Systems](#)
[Morpho Technologies](#)

[Motorola](#)
[Mysticom](#)
[Nazomi Communications](#)
[NEC Corp.](#)
[NewLogic](#)
[Nogatech](#)
[Nokia Mobile Phones](#)
[Nortel Networks](#)

[Nsine Limited](#)
[NurLogic Design](#)
[Nxtwave Communications](#)
[Oki Electric Industry](#)
[Palmchip](#)
[Philips Semiconductors](#)
[Piiri Technologies Oy](#)
[Pittsburgh Digital Greenhouse](#)

[Portability Inc.](#)
[Power X Limited](#)
[Qualis Design](#)
[RealChip](#)
[Robert Bosch GmbH](#)
[Samsung Electronics](#)
[Sanyo Electric](#)
[Seiko Epson](#)
[Semifore Technologies](#)

[Sharp](#)
[SIDSA](#)
[SI2](#)
[sci-Worx](#)
[Silicon Design Solutions](#)
[Silicon Metrics](#)
[Simplex Solutions](#)
[SIPAC](#)
[Sirius Communications NV](#)
[Siroyan Limited](#)
[Sonics Inc.](#)
[Sony](#)
[SPaSE BV](#)
[Spiratech Ltd](#)
[Spirea AB](#)

[STARC](#)
[ST Microelectronics](#)
[Synchronicity, Inc.](#)
[Synopsis](#)
[SynTest Technologies](#)
[Taiwan SoC Consortium](#)
[Telecom Italia Lab S.p.A](#)
[telefonaktiebolaget LM Ericsson](#)
[Teleraty Systems](#)
[Tensilica](#)
[Teradyne](#)
[Thales Group](#)
[The Silicon Group](#)
[THOMSON multimedia](#)
[Toshiba](#)
[TransEDA](#)
[TriMedia Technologies](#)

[TSMC](#)
[United Microelectronics Corp.](#)
[Vector 12](#)
[Verisity Design](#)
[Verplex Systems](#)
[Virtio](#)
[Virtual Component Exchange \(VCX\)](#)
[Virtual Silicon Technology](#)
[Vulcan Machines Limited](#)
[Xentec](#)
[Xi'an Swip Co. Ltd](#)
[Xilinx](#)
[X-VEIN](#)
[Xylon](#)

Individual Members

D.K. Arvind
Patrick Beauvillard
Annette Bunker
Guy Bois
Mark Buckner
Annette Bunker
Steven Burchfield

Magie Ge
Prof. Ganesh Gopalakrishnan
Dr. Peter Green
David Greenstein
Ken Hodor
Gerald Keeler
Edward Lee

Zainalabedin Navabi
Timothy O'Donnell
Brahmaji Potu, Ph.D.
Miodrag Potkonjak
Gang Qu
Wolfram Putzke Roeming
Eung Shin



Raymond Burkley Jean-Paul Calvez Eduardo Charbon Jong Kwan Choi Pai ChouKi-Seok Chung Lee Dilley	Kun-Bin Lee Tsung Lee Ian Mackintosh Samy Makar J. Sukarno Mertoguno	Angela Sutton Alberto Sangiovanni-Vincentelli Prab Varma Glenn Vinogradov Mason Weems
---	--	---

Last Edit - January 22, 2002

Copyright © 2002 VSI Alliance (tm)

<https://web.archive.org/web/20091024040005/http://www.guru3d.com/article/radeon-hd-5750-review-crossfirex/3>



You are here: Home » Hardware reviews » Videocards

Radeon HD 5750 review (Crossfire)

By: Hilbert Hagedoorn | Edited by John A. Johnsen | Published: October 21, 2009

Select Page to Load

Radeon HD Series 5700 features

So we have just established product positioning and have seen what cards to expect. Today is all about the Radeon HD 5750. Let's have a peek at some of the key features for this product:

- 1GB GDDR5 memory (128-bit)
- ATI Eyefinity technology with support for up to six displays
- ATI Stream technology
- Designed for DirectCompute 5.0 and OpenCL
- Accelerated Video Transcoding (AVT)
- Compliant with DirectX 11 and earlier revisions
- Supports OpenGL 3.1
- ATI CrossFireX™ multi-GPU support for highly scalable performance
- ATI Avivo™ HD video and display technology
- Dynamic power management with ATI PowerPlay technology
- DL-DVI, DVI, DisplayPort, HDMI
- PCI Express® 2.0 support

We'll address the majority of features in our article. But first let's focus on the sheer technical specifications. Transistor count for example. The number of transistors always works as an indicator of how powerful a product will be.

Today's tested product will position itself at the Radeon HD 4850 performance. The Radeon HD 4850 which we all know and love for its performance had 956 million transistors embedded onto that die. The new Radeon HD 5700 GPUs have 1040 million transistors. Correct, that is 1+ billion transistors tucked away in a small chip. The fabrication node just like the 5800 series is 40nm for this product.

The Radeon 5750 will be clocked at 700 MHz on the Core and Shader domain and 1150 MHz on the memory (4600 MHz on the memory), with the 5750 having 720 Shader processors. Though that looks a little pale in comparison to the 5800 series, remember .. these are mid-range products at really affordable prices.

The number of ROP units are rocking steady at 16 and sure -- texture units remain at 40 for the 5770 and 36 for the 5750.

But before you get blinded by all the specs in a few lines of text, let's break down the two cards announced today.

	Radeon HD 4870	Radeon HD 5750	Radeon HD 5770	Radeon HD 5870
Process	55nm	40nm	40nm	40nm
Transistors	956M	1.04B	1.04B	2.15B
Die Size	263 mm²	TBA	TBA	334 mm²
Core Clock	750 MHz	700 MHz	850 MHz	850 MHz
Shader Processors	800	720	800	1600
Compute Performance	1.2 TFLOPs	1.008 TFLOPs	1.36 TFLOPs	2.72 TFLOPs
Texture Units	40	36	40	80
Texture Fillrate	30.0 GTexels/s	25.2 GTexels/s	34 GTexels/s	68.0 GTexels/s
ROPs	16	16	16	32
Pixel Fillrate	12.0 GPixels/s	11.2 GPixels/s	13.6 GPixels/s	27.2 GPixels/s
Z/Stencil	48.0 GSamples/s	44.8 GSamples/s	54.4 GSamples/s	108.8 GSamples/s
Memory Type	GDDR5	GDDR5	GDDR5	GDDR5
Memory Clock	900 MHz	1150 MHz	1200	1200 MHz
Memory Data Rate	3.6 Gbps	4.6 Gbps	4.8 Gbps	4.8 Gbps
Memory Bandwidth	115.2 GB/s	73.6 GB/s	76.8 GB/s	153.6GB/s
Maximum Board Power (TDP)	160W	86W	108W	188W
Idle Board Power	90W	16W	18	27W

These numbers are really good, for the money this is going to make an excellent mid-range product series. We have not discussed it just yet, but memory; ATI will stick to DDR5 for both products. On the 5700 they'll crank it down a notch though we have 128-bit memory, cutting the bandwidth in half. However, since it's gDDR5 memory it will still offer sufficient enough bandwidth.

But yes -- the culprit of the 5700 series will be a cut off memory bandwidth, and this is the reason why it's performance will be slightly lower (on average). However, you can expect the Radeon HD 5750 to match any current single-GPU based graphics card in the mid-range segment like the Radeon HD 4850 and GeForce GTS 250. And all that with a single chip utilizing less than 86 Watts.

Power Consumption

One of the biggest accomplishments of the series 5000 graphics cards is the enhancement in the power design, the implementation of voltage and clock regulation is even more dynamic -- power management at a new level.

Site Navigation

- Home
- Latest News
- Submit News
- Hardware Reviews
- Articles & Guides
- VGA Charts new
- Game Reviews
- Forums
- Download Section
- Guru3D Price Grabber
- Guru Price Grabber UK
- Guru PC Buyers Guide
- Guru3D Stereo Section
- Guru3D Clan
- Guru3D Folding@Home
- Contact us
- Join our news-letter
- Set as Homepage

So we'll look purely at the Radeon HD 5770 now, in IDLE the GPU will clock down and lower its voltages on both GPU and memory. Have a look:

GPU	Radeon HD 4870	Radeon HD 5750	Radeon HD 5770
Max. Board Power (TDP)	160W	86W	108W
Idle Board Power	90W	16W	18W

The 5750 achieves a low 16W IDLE power consumption by clocking down in several power stages. Thus a low engine (core) clock frequency with lowered voltages and lower GDDR5 memory power. It's amazing though as your generic graphics card would normally consume 50~60 Watts when it idles in Windows.

Things get even better though, the performance of the graphics card opposed to the last generation products remains roughly the same yet the 5750 has a TDP (peak wattage) of only 86 Watts. That is just really good. We will test power consumption later on in this article.

Other Related Articles

- Radeon HD 5870 Review
- ATI Eyefinity review
- Radeon HD 5850 review (CrossfireX)
- Radeon HD 5770 review
- Radeon HD 5770 in 3-way CrossfireX review

Pages (24): « previous 1 2 [3] 4 5 next » ... Last »



Previous



Home

Check lowest prices on these products in **Guru3D.com price guide**, among the available categories: Retail & OEM Processors - Video Cards - Motherboards - Memory - Soundcards - Hard Drives - Monitors - Printers - DVDs - CD-RWs - PDAs and more !

Copyright (c) 1997-2009 Hilbert Hagedoorn, All Rights Reserved. Webdesign by Mohsin Ali - Legal disclaimer/notice
The Guru of 3D, the Hardware guru, and 3D Guru are the trademark ownership of Hilbert Hagedoorn.

⌵ Affiliates

- RivaTuner
- nVHardPage
- 3DMark Vantage
- SiSoft SANDRA
- Guru3D Driver Sweeper
- nVTempLogger
- ATI Tray Tools
- Reader Rig of the Month



⌵ Links

- Driver Scan
- Your company ?
- Your company ?

⌵ Downloads

- NVIDIA GeForce drivers
- ATI Catalyst drivers
- Benchmarks & Demo's
- Game Demo's
- NVIDIA Chipset drivers
- Intel Chipset drivers

⌵ Hot Downloads

- 3DMark Vantage
- SiSoft Sandra
- 3DMark06
- RivaTuner
- Crysis Benchmark
- nVHardPage



<https://web.archive.org/web/20090926023317/http://www.guru3d.com/article/radeon-hd-5870-review-test/3>



Adobe Flash Player is no longer supported

You are here: Home » Hardware reviews » Videocards

Radeon HD 5870 review

By: Hilbert Hagedoorn | Edited by Ian R. Barling & John A. Johnsen | Published: September 23, 2009

Select Page to Load

Radeon HD Series 5000 features

So we have just established product positioning and have seen what cards to expect. Today is all about the Radeon HD 5870. Let's have a peek at some of the key features for this product:

- 1GB GDDR5 memory
- ATI Eyefinity technology with support for up to six displays
- ATI Stream technology
- Designed for DirectCompute 5.0 and OpenCL
- Accelerated Video Transcoding (AVT)
- Compliant with DirectX 11 and earlier revisions
- Supports OpenGL 3.1
- ATI CrossFireX™ multi-GPU support for highly scalable performance
- ATI Avivo™ HD video and display technology
- Dynamic power management with ATI PowerPlay technology
- DL-DVI, DL-DVI, DisplayPort, HDMI
- PCI Express® 2.0 support

We'll address the majority of features in our article. But first let's focus on the sheer technical specifications. Transistor count for example. The number of transistors always works as an indicator of how powerful a product will be. For example, the Radeon HD 4870 which we all know and love for its performance had 956 million transistors embedded onto that die. The new Radeon HD 5800 GPUs have 2.15 billion transistors. Correct, that is 2150 million transistors tucked away in a small chip. The fabrication node was 40nm for this product, resulting in a die size of 334 mm², which for AMD is monolithic, yet thanks to the 40nm fabrication process only a third bigger than the previous 4890 GPU.

Now you'd think with so many transistors high clock frequencies would be an issue. Incorrect, the high-end Radeon HD 5870 will be clocked at a cool 850 MHz. In fact we even took it over 900 MHz without any issues.

The lower specced Radeon 5850 will be clocked at 725 MHz on the Core and Shader domain.

Shader processors then, we went from 800 Shader processors on the Radeon HD 4850/4870/4890 to 1600 shader processors (also called stream processors) on the Radeon HD 5870. That's doubled up. The ROPs went up from 16 to 32 as well and sure... texture units from 40 to 80 as well.

The Radeon HD 5850 will have some units cut away though. 1440 Shader cores, still we find 32 ROPS and 72 texture units. This product will for example be a good heap faster than say, a GeForce GTX 285.

But before you get blinded by all the specs in a few lines of text, let's break down the two cards announced today in comparison to last year's Radeon HD 4870.

	Radeon HD 4870	Radeon HD 5850	Radeon HD 5870
Process	55nm	40nm	40nm
Transistors	956M	2.15B	2.15B
Die Size	263 mm ²	334 mm ²	334 mm ²
Core Clock	750 MHz	725 MHz	850 MHz
Shader Processors	800	1440	1600
Compute Performance	1.2 TFLOPs	2.09 TFLOPs	2.72 TFLOPs
Texture Units	40	72	80
Texture Fillrate	30.0 GTexels/s	52.2 GTexels/s	68.0 GTexels/s
ROPs	16	32	32
Pixel Fillrate	12.0 GPixels/s	23.2 GPixels/s	27.2 GPixels/s
Z/Stencil	48.0 GSamples/s	92.8 GSamples/s	108.8 GSamples/s
Memory Type	GDDR5	GDDR5	GDDR5
Memory Clock	900 MHz	1000 MHz	1200 MHz
Memory Data Rate	3.6 Gbps	4.0 Gbps	4.8 Gbps
Memory Bandwidth	115.2 GB/s	128.0 GB/s	153.6GB/s
Maximum Board Power (TDP)	160W	170W	188W
Idle Board Power	90W	27W	27W

These numbers are downright staggering. We have not discussed it just yet, but memory; ATI will stick to DDR5 for both products. On the 5870 they'll up it a notch alright as the clock frequency has been upped a notch as well, as some design changed in the memory controllers.

Expect the Radeon HD 5870 to outperform any current single-GPU based graphics card like the Radeon HD 4890 and GeForce GTX 285. And all that with a single chip utilizing less than 190 Watts.

Radeon HD 5870

Site Navigation

- Home
- Latest News
- Submit News
- Hardware Reviews
- Articles & Guides
- VGA Charts [new](#)
- Game Reviews
- Forums
- Download Section
- Guru3D Price Grabber
- Guru Price Grabber UK
- Guru PC Buyers Guide
- Guru3D Stereo Section
- Guru3D Clan
- Guru3D Folding@Home
- Contact us
- Join our news-letter
- Set as Homepage



Adobe Flash Player is no longer supported

Other Related Articles

- AMD ATI Radeon 3850 & 3870 review
- Silicon Power M10 External SSD review
- Inno3D GTS 250 TwinTurbo Pro review
- MSI Radeon 4890 Cyclone SOC review (1GHz)
- Radeon HD 5870 CrossfireX test review

Pages (27): « previous 1 2 [3] 4 5 next » ... Last »



Previous



Home

Check lowest prices on these products in **Guru3D.com price guide**, among the available categories: Retail & OEM Processors - Video Cards - Motherboards - Memory - Soundcards - Hard Drives - Monitors - Printers - DVDs - CD-RWs - PDAs and more !

Copyright (c) 1997-2009 Hilbert Hagedoom, All Rights Reserved. Webdesign by Mohsin Ali - Legal disclaimer/notice
The Guru of 3D, the Hardware guru, and 3D Guru are the trademark ownership of Hilbert Hagedoom.

(v) Affiliates

- RivaTuner
- nVHardPage
- 3DMark Vantage
- SiSoft SANDRA
- Guru3D Driver Sweeper
- nVTempLogger
- ATI Tray Tools
- Reader Rig of the Month



(v) Links

- Driver Scan
- Your company ?
- Your company ?

(v) Downloads

- NVIDIA GeForce drivers
- ATI Catalyst drivers
- Benchmarks & Demo's
- Game Demo's
- NVIDIA Chipset drivers
- Intel Chipset drivers

(v) Hot Downloads

- 3DMark Vantage
- SiSoft Sandra
- 3DMark06
- RivaTuner
- Crysis Benchmark
- nVHardPage



AMD Evergreen: 40nm DirectX 11 GPU Silicon Arrives



Date: June 3rd, 2009
Topic: Video Card
Manufacturer: ATI
Author: Derek Wilson

Buy the HIS H489F1GP Radeon HD Video Card PCI



\$218.99
20.00 rebate




\$219.99
20.00 rebate



\$224.99
20.00 rebate

Today marks our first glimpse at Evergreen.



Evergreen is the codename for AMD's 40nm DX11 based product. With AMD's codename shift this time around, it is still a little ambiguous as to whether Evergreen is a GPU class or a specific configuration. While it is exciting that they're sharing what they are, AMD are still holding quite a few cards close to their chest. We know almost nothing about the details of the configuration, except that it is built on a 40nm process, is fully DX11 compliant, and that AMD put forth a slide titled "DirectX 11: sooner than you think."

Of course, sooner than I think would have to mean DirectX 11 would be here before Windows 7 launches (which is sort of impossible). We learned today that Microsoft has announced Windows 7 will be here on shelves on October 22nd. From AMD's press release and slide deck, however, the only safe bet on when we'll see hardware is before the end of the year. But teasing us with "sooner than you think" would just be mean if we'll have to wait that long.

Latest news by DAILY TECH

June 5, 2009

FTC Closes Rogue ISP
Japan Offers Incentives to Purchase Green Products
Russia Attracts Another Space Tourist: Cirque du Soleil Owner
Apple's Steve Jobs Eyes Return to Chief Role
Intel Buys Wind River for \$884M
Alltel Cuts Contract Length in Half
Apple to Bring New Data Center to North Carolina

June 4, 2009

AT&T, Verizon, and Sprint Can't Be Sued for Domestic Spying
Game Industry Remains Optimistic Despite Economy
SanDisk CEO: "You Can't out-iPod the iPod"
Black Nintendo Wii Coming to Japan
Sony Says no PS3 Price Cuts Coming
First Official Palm Pre Reviews Hit the Net
NASA Study Acknowledges Solar Cycle, Not Man, Responsible for Past Warming
Microsoft CEO to Obama: Drop Tax Break Cuts or We'll Send Jobs Overseas
Scientists Create Flexible Memristor, Holds Promise of Flexible Memory
Researchers Restore Vision to Legally Blind With Stem Cells

[more Headlines](#)

AnandTech Blogs

Computex 2009 - Gigabyte
Computex 2009 - ASRock
Computex 2009 - Intel ULV Processor Updates

Latest Video Card Forum Posts

Need Upgrade Help
What is the Best LCD Under \$1000
RMA for fun and profit!
Video Card Recommendations for 30" Monitor
STICKY: The LCD Thread
Most powerful gfx card?
Corsair HX620 handle GTX 280 SLI and Core i7 920?
Acer x223w wide mode?
I need a 22" monitor for gaming and movies
What happened to Lucid's Hydra chip?

[more Video Card Discussions](#)

If AMD doesn't need to respin their silicon and Evergreen really is complete as of today (aside from qualification, tuning and ramping production), AMD really could be ready to ship product much sooner than November, getting product out along side or in advance of Windows 7. It all really depends on a bunch of things they aren't going to tell us. Of course that's just the way it is, but that doesn't make it any less frustrating not knowing.

We can ascertain from the wafer shot AMD provided us that there are about 19.5 dies vertically and 25.5 dies horizontally. As this is a 300mm wafer, we can sort of "guess" the dimensions of the chip at roughly 15.38mm x 11.76mm resulting in a die area of a little over 180mm². This is much smaller than the 55nm RV770 which is a 260mm² part, so if we expect a similar price target for the first Evergreen die that we saw with RV770, we could see a significant cost savings (and hopefully this tiny die will deliver a good general performance improvement as well).

Compare this die size to RV740 that weighs in at 137mm² and 826 million transistors, and we can estimate very loosely that Evergreen could come in at something over 1 billion transistors. Certainly the process has been tweaked since RV740 and Evergreen is a different architecture and layout, so scaling isn't exact or direct, but proportionally AMD should be able to fit that many transistors into only 180mm² on the 40nm process.

We'll definitely be interested in seeing how close our guesses are when we finally have hardware in hand. But this still doesn't speak to performance. Our only real reference for AMD's target is to guess that they will want to come in at near the same price point the RV770 hit at: the \$200 - \$300 range. With the added transistors, changes in architecture, and clock speed advantages with the 40nm process over 55nm, we could reasonably see AMD hitting that target.

Showing off Demos and Preliminary Thoughts

Unfortunately, we don't have any videos of the demos in action as presented at Computex. All we have to go on right now is a presentation from AMD with some slides showing stills from some demos. Here's a DX11 SDK sample showing sub-division surfaces using tessellation.



Since we don't have hardware that can run this demo as of yet, we can't really know how good 14 frames per second really is or what it means. We'll have to wait until we have competitive hardware to do any evaluation of performance (or we'll have to wait until AMD shows us something we can compare to current DX10 hardware running on Evergreen). Also shown were some other demos including transcoding and compute shader based AI.



While there are a lot of applications that will be potentially enabled by the compute shader, it will certainly take some time for developers to get familiar and comfortable with treating the GPU like another processor for data-parallel compute operations. Our expectation is that DX11 (and preparation for the next round of console game systems) will serve as a catalyst for a final push away from DX9 based hardware and into the realm of real programmability.

While dynamic branching has been available since DX9, we haven't seen heavy use of it as it can be very resource intensive and slow even on modern hardware. Yes, it's much faster today, but there are still some things that are not practical today despite the fact that some algorithms would significantly benefit from fast branching. We hope, and it seems AMD believes, that DX11 class hardware will continue the trend of speeding up and improving branching on GPUs. Here's a timeline showing ATI's GPU technology progression at a very high level.



From the context of the presentation and what we've seen, AMD is really interested in pushing their tessellation technology. Tessellation has been an option since R600, but it hasn't been taken incredibly seriously because it wasn't a standard feature and couldn't run on all DX10 hardware. DirectX 11 changes that, and AMD is touting their experience as a benefit here. We'll have to wait and see what NVIDIA does, as experience doesn't always mean better.

So, the big deal today is that AMD is showing off working DX11 silicon on a 40nm process running DX11 demos. We expect DX11 before the end of the year based on what AMD has said. Now that we know when Windows 7 will hit (October 22nd) and that AMD already has Evergreen silicon in their hands, we are fairly hopeful that AMD could introduce their DX11 part ahead of Windows 7 if everything goes smoothly for them from now until then. Definitely cool stuff that really puts the pressure on NVIDIA to follow suit and start talking about their DX11 answer.

[Home](#)

Index

Article Index

Tools

Find the lowest prices

Print this article

Share

Digg

del.icio.us

E-mail

31 Comments - Last by Obsy, 1 days ago

[Login](#) | [Create Account](#)

Did you say 1B transistors?? by [Dante80](#), 2 days ago

So, from the looks of it (I assume we are seeing RV870 and not RV840 there), RV870 would be a tweaked for DX11 RV790 @40nm, without an increase in shader count.

If AMD goes again for best value/cost for the performance segment (employing the same binning scalability that the RV770 strategy used), they will have the same competitive products in the 200-300\$ range, with better margins per chip (but less yields are implied, at least atm due to TSMC difficulties with the process).

Can that be enough to counter the Nvidia loch ness monster?...

Also, will it blend?...XD

ps: First post in AT :)

RE: Did you say 1B transistors?? by [samsprqr](#), 2 days ago

this new chip should have around 13% more transistors than RV790

if what charlie demerjian says is true, that most of DX11 features were already implemented in old ati chips, there could be some resource increases (more shaders)

yet, if that's not to come, and clock rates are not improved by much either because of tsmc's problems, it'd be hard to try and sell this stuff for \$300

RE: Did you say 1B transistors?? by [psychobriggsy](#), 2 days ago

Maybe ATI is going for clock speed increases over shader increases.

This Evergreen might have, say, 960 shaders as opposed to the 800 in RV790 (I'm sure someone can do a 40nm die analysis to work out what percentage of the RV740 die 640 shaders are, then work out how many there would be on the bigger RV870 die, assuming non-shader stuff was pretty much the same).

But maybe the shader clocks are higher (as with NVIDIA's offerings), so AMD can offer good performance with small (and hence cheap) die sizes.

Looking Forward to it. by [Iwodo](#), 2 days ago

Looks like it will be MUCH more profitable then Rv770. Yield problem as mentioned above would be improved by the time they ship the first batch.

However i have doubt on performance improvement. Consider the extra transistors are needed for Direct X Compute and Direct X 11 features. Apart from Clock speed i dont see any speed improvement over Rv870. (I hope i am wrong)

I hope we get UVD improvement as well.

RE: Looking Forward to it. by [plonk420](#), 2 days ago

I'm more interested in GPGPU/DX11/OpenCL capability ... and just a WEEEE bit of gaming ;)

About Nvidia. by [Iwodo](#), 2 days ago

I forgot Nvidia as well. Since my beloved Apple is in bed with Nvidia. I seriously hope Nvidia has something good to counter it. On paper Greenland looks like to offer some very good performance / watt / price....

Dual core GPU? by JAKra, 2 days ago

The core looks so symmetrical, like a dual core CPU.
Anyway it looks way different than RV770.
180mm^2 for a performance part? I'm not used to this... :D
Eager to see the performance. :)

So how 'bout a public demo release this time? by seriouscat, 2 days ago

After all the promises of Cinema 2.0 and then having NONE of the demos being released to the public, I'll sit back and wait on this one.

More Thoughts on ATI by lwodo, 2 days ago

Arh, i forgot to mention how Nvidia did so well when ATI were split into three group to design chips for Xbox 360, Wii and the PC part. After everything was finished ATI seems to hit back very hard. Any thoughts?

Waiting for DX12 by MadBoris, 2 days ago

DX9 is still here to stay and will be more popular until the next console refresh.

Unfortunately consoles drive the state of hardware adoption and game integration, DX11 still won't be really utilized or leveraged until the next gen console comes out.

Glad to see technology still moving forward but it must be hard to produce technology and hardware knowing it's only a building block towards something else in the future when the next gen consoles set a new hardware standard.

New GPU's need to provide considerably more performance with enhanced features (not DX11 features that devs won't use) then I may consider a new GPU.
My 8800 still plays all multiplatform games fine so finding a good justification for upgrading is still difficult.

I wonder how long it will take until these GPU companies realize they can't spit out hardware like this without losing money when developers are now tuned to console product cycles.

Comments Page 1 of 4

[1] [2] [3] [4] Next

Sponsored Links (Get Listed)

[Try Symantec Protection Suite today](#)

Protect Your Information Against More Risks In More Places.

[Know a company guilty of software piracy?](#)

Get PAID up to \$1 Million for Confidentially Reporting Software Piracy by an organization.

[Stop buying storage](#)

Start Optimizing Data. Calculate Your Storage Savings Today.

[HP Promotions & Rebates for Desktops](#)

New SYNEX exclusive discounts on select HP desktop parts.

[Vipower DVD Duplicator](#)

Manufacturer of quality disc duplication systems. Drop shipping available. Become a reseller today.

Intel's 32nm Update: The Follow-on to Core i7 and More

Date: Feb 11, 2009
Type: CPU & Chipset
Manufacturer: [Intel](#)
Author: [Anand Lal Shimpi](#)
Page 1

Seven billion dollars.
That's the amount that Intel is going to spend in the US alone on bringing up its 32nm manufacturing process in 2009 and 2010.
These are the fabs Intel is converting to 32nm:

32nm Manufacturing Fabs



In Oregon Intel has the D1D fab which is already producing 32nm parts, and D1C which is scheduled to start 32nm production at the end of this year. Then two fabs in Arizona: Fab 32 and Fab 11X. Both of them come on line in 2010.
By the end of next year the total investment just to enable 32nm production in the US will be approximately eight billion dollars. In a time where all we hear about are bailouts, cutbacks and recession, this is welcome news.
If anything, Intel should have a renewed focus on competition given that its chief competitor finally woke up. That focus is there. The show must go on. 32nm will happen this year. Let's talk about how.

Page 2
The Manufacturing Roadmap


The tick-tock cadence may have come about at the microprocessor level, but its roots have always been in manufacturing. As long as I've been running AnandTech, Intel has introduced a new manufacturing process every two years. In fact, since 1989 Intel has kept up this two year cycle.

Process Name	P1264	P1266	P1268	P1270	P1272
Lithography	65nm	45nm	32nm	22nm	16nm
1 st Production	2005	2007	2009	2011	2013

Manufacturing

Development

We saw the first 45nm CPUs with the Penryn core back in late 2007. Penryn, released at the very high end, spent most of 2008 making its way mainstream. Now you can buy a 45nm Penryn CPU for less than \$100.
The next process technology, which Intel refers to internally as P1268, shrinks transistor feature size down to 32nm. The table above shows you that first production will be in 2009 and, after a brief pause to check your calendars, that means this year. More specifically, Q4 of this year.
I'll get to the products in a moment, but first let's talk about the manufacturing process itself.
Here we have our basic CMOS transistor:




Current flows from source to drain when the transistor is on, and it isn't supposed to flow when it's off. Now as you shrink the transistor, all of its parts shrink. At 65nm Intel found that it couldn't shrink the gate dielectric any more without leaking too much current through the gate itself. Back then the gate dielectric was 1.2nm thick (about the thickness of 5 atoms), but at 45nm Intel's switched from a SiO₂ gate dielectric to a high-k one using Hafnium. That's where the high-k comes from.

The gate electrode also got replaced at 45nm with a metal to help increase drive current (more current flows when you want it to). That's where the metal gate comes from.

The combination of the two changes to the basic transistor gave us [Intel's high-k + metal gate transistors at 45nm](#), and at 32nm we have the second generation of those improvements.

The high-k gate dielectric gets a little thinner (equivalent to a 0.9nm SiO₂ gate, but presumably thicker since it's Hafnium based, down from 1.0nm at 45nm) and we've still got a metal gate.



At 32nm the transistors are approximately 70% the size of Intel's 45nm hk + mg transistors, allowing Intel to pack more in a smaller area.

- 2nd generation high-k + metal gate transistors
- Immersion lithography on critical layers
- 9 copper + low-k interconnect layers
- ~70% dimension scaling from 45nm generation
- Pb-free and halogen-free packages


The big change here is that Intel is using immersion lithography on critical metal layers in order to continue to use existing 193nm lithography equipment. The smaller your transistors are, the higher resolution your equipment has to be in order to actually build them. Immersion lithography is used to increase the resolution of existing lithography equipment without requiring new technologies. It is a costlier approach, but one that becomes necessary as you scale below 45nm. Note that [AMD to made the switch to immersion lithography at 45nm](#).

- 2nd generation high-k + metal gate
 - 0.9nm equivalent oxide thickness high-k (scaled from 1.0 nm on 45nm)
 - Replacement Metal Gate process flow
 - 30nm gate length
 - 4th generation strained silicon
- >22% performance increase
- Tightest reported gate pitch

- Highest reported drive currents

Intel reported significant gains in transistor performance at 32nm; the graphs below help explain:

32nm Transistor Performance



32nm provides improved performance or reduced leakage
Highest drive current of all reported 32nm technologies

We're looking at the comparison of leakage current vs. drive current for both 32nm NMOS and PMOS transistors. The new transistors showcase a huge improvement in power efficiency. You can either run them faster or run them at the same speed and significantly reduce leakage current by a magnitude of greater than 5 - 10x compared to Intel's 45nm transistors. Intel claims that its 32nm transistors boast the highest drive current of all reported 32nm technologies at this point, which admittedly there aren't many.


The power/performance characteristics of Intel's 32nm process make it particularly attractive for mobile applications. But more on that later.

Page 3

Fat Pockets, Dense Cache, Bad Pun

Whenever Intel introduces a new manufacturing process the first thing we see it used on is a big chip of cache. The good ol' SRAM test vehicle is a great way to iron out early bugs in the manufacturing process and at the end of 2007 Intel demonstrated its first 32nm SRAM test chip.


- 0.171 μm^2 cell
- 291 Mbit
- >1.9 billion transistors
- 4 GHz operation
- First demonstrated Sep '07



Intel's 32nm SRAM test vehicle


The 291Mbit chip was made up of over 1.9 billion transistors, switching at 4GHz, using Intel's 32nm process. The important number to look at is the cell size, which is the physical area a single bit of cache will occupy. At 45nm that cell size was 0.346 μm^2 (for desktop processors, Atom uses a slightly larger cell), compared to 0.370 μm^2 for AMD's 45nm SRAM cell size. At 32nm you can cut the area nearly in half down to 0.171 μm^2 for a 6T SRAM cell. This means that in the same die area Intel can fit twice the cache, or the same amount of cache in half the area. Given that Core i7 is a fairly large chip at 263 mm^2 I'd expect Intel to take the die size savings and run with them. Perhaps with a modest increase to L3 cache size.

SRAM Cell Size Scaling



A big reason we're even getting this disclosure today is because of how healthy the 32nm process is. Below we have a graph of defect density (number of defects

One big reason we're even getting this disclosure today is because of how leading the 32nm process is. Below we have a graph of defect density (number of defects in silicon per area) vs time; manufacturing can't start until you're at the lowest part of that graph - the tail that starts to flatten out:



Intel's 45nm process ramped and matured very well as you can see from the chart. The 45nm process reached lower defect densities than both 65nm and 90nm and did it faster than either process. Intel's 32nm process is on track to outperform even that.

Two Different 32nm Processes?

With Intel now getting into the SoC business (System on a Chip), each process node will now have two derivatives - one for CPUs and one for SoCs. This started at 45nm with process P1266.8, used for Intel's consumer electronics and Moorestown CPUs and will continue at 32nm with the P1269 process.

	45 nm		32 nm		22 nm	
Process:	P1266	P1266.8	P1268	P1269	P1270	P1271
Products:	CPU	SoC	CPU	SoC	CPU	SoC

There are two major differences between the CPU and SoC versions of a given manufacturing process. One, the SoC version will be optimized for low leakage while the CPU version will be optimized for high drive current. Remember that graph comparing leakage vs. drive current of 45nm vs. 32nm? The P1268 process will exploit the arrows to the right, while P1269 will attempt to push leakage current down.


The second difference is that certain SoC circuits require higher than normal voltages and thus you need a process that can tolerate those voltages. Remember that with a SoC it's not always just Intel IP being used, there are many third parties that will contribute to the chips that eventually make their way into smartphones and other ultra portable devices.

The buck doesn't stop here, in two more years we'll see the introduction of P1270, Intel's 22nm process. But before we get there, there's a little stop called Sandy Bridge. Let's talk about microprocessors for a bit now shall we?

Page 4

Tick-Tock: U R Doin it Right

Let's check the stats; Conroe in July 2006, Penryn in October 2007, Nehalem in November 2008. That's a tock, tick, and another tock, each about a year apart. Note that the cadence does appear to be slipping a bit, but we'll see exactly when in 2009 we get Westmere before making any accusations.



The next tick is, as I just mentioned, Westmere. It's a 32nm shrink of Nehalem, much like Penryn was a 45nm shrink of Conroe/Merom. And it's due out in the fourth quarter of this year.

Yesterday, Intel demonstrated working versions of its 32nm processors in both desktops and notebooks. The notebook aspect of the demonstration is very important, which I'll get to later. Both mobile and desktop versions of Westmere will be shipping from Intel in Q4.

Getting Complicated with Code Names

Nehalem is the overall name for Intel's 45nm desktop/mobile/server product family. At the high end we have Bloomfield, which is the quad-core, eight-thread, Core i7 processor we all long for. That's the only Nehalem derivative that's launched thus far.

Segment	Manufacturing Process	Socket	Processor	Cores	Threads	Release Date
High End	45nm	LGA-1366	Bloomfield	4	8	Q4 2008

Desktop		1300				2008
Mainstream Desktop	45nm	LGA-1156	Lynnfield	4	8	2H 2009
Mobile	45nm	mPGA-989	Clarksfield	4	8	2H 2009
4S Server	45nm	LGA-1567	Nehalem-EX	8	16	2H 2009
2S Server	45nm	LGA-1366	Nehalem-EP	4	8	1H 2009
1S Server	45nm	LGA-1156	Lynnfield	4	8	2H 2009

By the end of this year we'll see Lynnfield and Clarksfield. These are both quad-core, eight-thread Nehalem processors but at lower TDPs and price points. They will fit into Intel's unannounced LGA-1156 socket and only support two channels of DDR3 memory (compared to LGA-1366 and 3-channels with Core i7).

On the server side we'll have Nehalem-EX, an 8-core, 16-thread version. Nehalem EP a 4-core, 8-thread version. And Lynnfield again for the entry level servers.

These are all 45nm parts and all due out by the end of this year.

Note that there's one name missing: Havendale. Havendale was supposed to be a 2-core Lynnfield + on-chip graphics, perfect for notebooks and low end desktops where quad-core isn't necessary. Unfortunately, Havendale got delayed until Q4 2009 with systems shipping in Q1 2010. That just happened to coincide with Intel's 32nm ramp so a very significant decision was made: Havendale got scrapped.

Page 5

Enter the 32nm Lineup

Instead of Havendale in Q4, we'll get Clarkdale and Arrandale. These are both dual-core, quad-thread processors, and both have on-package graphics. The CPU cores will be built on Intel's 32nm process and in fact, they will be the first Westmere CPUs shipping into the market.




Now note that the dual-core market is the largest slice of the processor pie. Intel must be incredibly confident in its 32nm process to start shipping it into these demand markets first. Remember that both 65nm and 45nm initially launched on the high end desktop, but 32nm is making its debut in mainstream notebooks and desktops. The 32nm ramp is going to be a good one folks.

Segment	Manufacturing Process	Socket	Processor	Cores	Threads	Release Date
High End Desktop	32nm	LGA-1366	Gulftown	6	12	1H 2010
Mainstream Desktop	32nm	LGA-1156	Clarkdale	2	4	Q4 2009
Mobile	32nm	mPGA-989	Arrandale	2	4	Q4 2009
4S Server	32nm	LGA-1567	???	?	?	2010
2S Server	32nm	LGA-1366	???	?	?	2010
1S Server	32nm	LGA-1156	Clarkdale	2	4	2010

Clarkdale/Arrandale have 32nm CPUs but their on-package GPUs are still built on Intel's 45nm process; these are the GPUs that were supposed to be used for Havendale! It won't be until 2010 with Sandy Bridge that we see a 32nm CPU and 32nm GPU on the same package.





A side effect of the Clarkdale/Arrandale architecture is that the memory controller is now located on the GPU and not the CPU, although both are still on package and should still be quite low latency.

Westmere Family Mainstream Desktop / Mobile Processor

Key Features¹:

- Intel microarchitecture codename Nehalem on 32nm
- Multi-Chip Package (MCP) processor with:
 - processor cores built on 32nm 2nd generation high-k metal gate process
 - Integrated graphics controller & memory controller built on 45nm high-k metal gate process
- Intel Turbo Boost technology
- Intel Hyper-Threading technology (2 Cores, 4 threads)
- Intel Smart Cache
- Integrated memory controller (IMC)
- Integrated, discrete / switchable graphics support
- Advanced Encryption Standard (AES) acceleration
- Compatible with the Intel 5 series chipset based platforms due to be released in 2H09


Worlds First 32nm Based Processor

Keep following; if you want a quad-core Westmere, your only option will be in the LGA-1366 socket with Gulftown. **Core i7 will get replaced with a six-core, twelve-thread processor in early 2010.** There won't be a 32nm quad-core part on the desktop until the end of 2010 with Sandy Bridge.

Page 6

The Server Roadmap

Intel's 32nm server roadmap is notably different from the desktop roadmap. Nehalem-EX will ship into the Xeon 7000 series as an 8-core, 16-thread part. It will eventually get replaced sometime in 2010 with a 32nm Westmere derivative.




We'll see a 32nm six-core Westmere based processor in the Xeon 5000 series in 2010.

Finally Lynnfield and Clarkdale will be carried over to the entry level Xeon platforms at the end of this year and into 2010.


What About Chipsets?

Intel's X58 chipset will remain the top dog through 2010. Chances are that we won't see it replaced until the next tock with Sandy Bridge. Now that isn't to say that the six-core 32nm Gulftown will work in existing X58 motherboards; while that would be nice, Intel does have a habit of forcing motherboard upgrades, we'll have to wait and see.

The rest of the Nehalem/Westmere family will rely on Intel's upcoming P55 chipset:



Originally both Lynnfield and Havendale were to have an on-package PCIe controller, I'm not sure if that has changed with the Havendale cancellation but I see no reason for it to have. In which case a Lynnfield system will still look like this:



Page 7

Westmere's New Instructions

Much like Penryn and its new SSE4.1 instructions, Westmere comes with 7 new instructions added to those already in Core i7. These instructions are specifically focused on accelerating encryption/decryption algorithms. There's a single carryless multiply instruction (PCLMULQDQ...I love typing that) and 6 instructions of AES.

Westmere: AES New Instructions

- Use additional transistor budget to add new capabilities
 - Similar to adding SSE4.1 in Penryn (45nm tick)
- 7 new instructions for accelerating encryption/decryption algorithms
 - Carryless multiply (PCLMULQDQ)
 - 6 instructions for AES
- Example client usage
 - Enables full disk encryption

Intel gives the example of hardware accelerated full disk encryption as a need for these instructions. With the new instructions being driven into the mainstream

first, we'll probably see quicker than usual software adoption.

Final Words

What is there to say other than: it's a healthy roadmap. The only casualty I've seen is Havendale but I'd gladly trade Havendale for a 32nm version. But let's get down to what this means for what you should buy and when.

At the very high end, Core i7 users have little reason to worry. While Intel is expected to bump i7 up to 3.33GHz in the near future, nothing below i7 looks threatening in 2009. Moving into 2010, the 6-core 32nm i7 successor should be extremely powerful. Intel's strategy with LGA-1366 makes a lot of sense: if you want more cores, this is the platform you're going to have to be on.

Now although I said that nothing will threaten Core i7 this year, you may be able to get i7-like performance out of Lynnfield in the second half. A quad-core Lynnfield running near 3GHz, should offer much of the performance of an i7 with a lower platform cost. Remember back to our original i7 review; we didn't find a big performance benefit from three channels of DDR3 versus two.

Lynnfield is on track for a 2H 2009 introduction and if you're unable to make the jump to i7 now, you'll probably be able to get i7-like performance out of Lynnfield in about 6 months. Intel did mention that the most overclockable processors would go into the LGA-1366 socket. Combine better overclockability with the promise of 6 cores in the future and it seems like LGA-1366 is shaping up to be a platform that's going to stick around despite cheaper alternatives.

The 32nm Clarkdale/Arrandale parts arriving by the end of this year really means one very important thing: the time to buy a new notebook will be either in Q4 2009 or Q1 2010. A 2-core, 4-thread 32nm Westmere derivative is not only going to put current Penryn cores to shame, it's going to be extremely power efficient. In its briefing yesterday, Intel mentioned that while Clarkdale/Arrandale clock speeds and TDPs would be similar to what we have today, you'll be getting much more performance. Seeing what we've seen thus far with Nehalem, I'd say a 2-core, 32nm version in a notebook is going to be reason enough for you to want to upgrade.

If I had to build a new desktop today I'd go Core i7 and think about upgrading to a 6-core version sometime next year. If I couldn't or didn't need to build today, then the thing to wait for is Lynnfield. Four cores that should deliver i7-like performance just can't be beat, and platform costs will be much cheaper by then (expect ~\$100 motherboards and near price parity between DDR3 and DDR2).

On the mainstream quad-core side, it may not make sense to try to upgrade to 32nm quad-core until Sandy Bridge at the end of 2010. If you buy Lynnfield this year, chances are that you won't feel a need to upgrade until late 2010/2011.

On the notebook side, if I needed one today I'd buy whatever I could keeping in mind that within a year I'm going to want to upgrade. I mentioned this in much of my recent Mac coverage; if you bought a new MacBook, it looks great, but the one you're going to really want will be here in about a year.

We owe Intel a huge thanks for being so forthcoming with its roadmaps. It's going to be a good couple of years for performance.

Intel Core 2 Extreme QX9650 - Penryn Ticks Ahead

Date: Oct 29, 2007

Type: CPU & Chipset

Manufacturer: [Intel](#)

Author: [Anand Lal Shimpi](#)

Page 1

We've never seen Intel with such a strong roadmap before, the company is truly firing on all cylinders and executing with amazing precision. Server, mobile, desktop and even new areas like ultra mobility and graphics all have absolutely wonderful roadmaps to look forward to. The biggest complaint we've had about Intel these days is that they kind of botched the X38 launch. Think back to a couple of years ago, what was our chief complaint then? Probably leaving us with power hungry, under performing processors for about 5 years. Today we're looking at a damn good Intel.


Recap: What's a Penryn?

Core, it's the architecture that shook an industry and today Intel is officially doing its first update to it. Prior to Intel's Core architecture, there wasn't much to get excited about when it came to Intel on the desktop.

At the same time, with Penryn Intel is very much the victim of its own success. How do you follow up such a tremendous splash with anything but equal greatness? AMD is close but still has yet to produce a response to Core 2, much less Penryn, and thus Intel's biggest competition today is itself.

In January 2007 Intel first showed off its 45nm High K + metal gate transistors, a dramatic departure from Intel's current 65nm transistors not only in size/switching speed but actual composition. If you remember back to the days of the original P6 processors, with a smaller transistor we saw tremendous improvements in die size, power consumption and performance. These days, such dramatic improvements are much harder to come by given that we're already dealing with such small transistor feature sizes. Gone are the days of the free lunch with each die shrink.

We've gone over [the technical details of Intel's high-K + metal gate enhancements](#), but the end result is that at the same clock speed you can expect dramatic reductions in power. Alternatively, at the same power levels, you can achieve much higher switching rates and thus higher clock speeds.



The core architecture of Penryn remains unchanged from Conroe; with the smaller transistors Intel's able to fit in a few new features and more cache on the chip while still maintaining a smaller die size. Where each dual-core Conroe die measured 143 mm², Penryn is merely 107 mm² despite having 50% more cache (6MB vs. 4MB). Obviously the quad core chips double overall area but you get the point.

Intel also uses a lot more of these new 45nm transistors than before; while a dual-core Conroe was made up of 291 million transistors, the comparable Penryn weighs in at 410 million (582M vs. 820M for quad-core variants). You're getting 40% more transistors and 50% more cache in a 25% smaller package; the latter is obviously most important to Intel as it helps reduce costs and drive profits up. So while it may seem generous, the move is purely self motivated on Intel's part.


The larger cache is a bit different than what we've seen in Conroe. While Conroe's cache is a 4MB 16-way set-associative L2, the 6MB Penryn cache is 24-way set-associative, designed to improve hit rates and keep latency manageable in an already large cache. Intel hasn't revealed whether Penryn's prefetchers have been adjusted to help populate its larger cache any better. As we saw in [our original Penryn preview](#), Penryn's cache performance remains unchanged; latencies in our final stepping are identical to Conroe.

The cache enhancements are by far the biggest consumer of those extra transistors in Penryn, but believe it or not, they aren't responsible for the biggest performance boost. Intel has been fairly steady in adding new instructions to the x86 ISA and Penryn continues the trend with the addition of SSE4. Penryn gets 47 new instructions that make up the first implementation of SSE4; more will come with Nehalem at the end of 2008. We'll talk about SSE4 performance later on in this article, but here are the instructions you get with Penryn:


Sub Group	Instructions	Description	Expected Application Benefits
Packed Dword Multiplies	PMULDQ, PMULBQ	New support for four signed or unsigned 32x32 bit multiplications per instruction as well as signed forms of 32x32->64 multiplication.	Broadly useful for improved automated compiler vectorization of data processing written in high-level languages like C and Fortran.
Floating Point Dot Product	DPBS, DPBSQ	Improved performance for AGS (Array of Structures) data processing through support for single and reusable precision dot products.	3D content creation, gaming, and support for languages like C++ and MATLAB.
Packed Blending	BLENDPS, BLENDQDQ, BLENDQPS, BLENDQDQ, BLENDQPS, BLENDQDQ	Blending conditionally copies one field in the source onto the same field in the destination. These new instructions improve the performance of blending operations for most field sizes through packing multiple operations in a single instruction.	Broadly useful for automated compiler vectorization of data processing written in high-level languages like C and Fortran, and applications such as image processing, video processing, multimedia, and gaming.
Packed Integer Min and Max	PMINSD, PMAXSD, PMINQD, PMAXQD, PMINQD, PMAXQD, PMINQD, PMAXQD	Compares packed signed/unsigned bytes/words/dword integers in the destination operand and the source operand, and returns the minimum or maximum value.	Broadly useful for automated compiler vectorization of data processing written in high-level languages like C and Fortran, and applications such as image processing, video processing, multimedia, and gaming.

		These new instructions support for vector packed operands in the destination operands.	These new instructions support for vector packed operands in the destination operands.
Floating Point Round	ROUNDPS, ROUNDSS, ROUNDPD, ROUNDDQ	Efficiently rounds the scalar and packed single- and double-precision operands to integers, with enhanced support for Fortran, Java and C/C++ language requirements.	Image processing, graphics, video processing, 2-D/3-D applications, multimedia, and gaming.
Register Insertion/Extraction	INSERTPS, PSHRQB, PSHRQB, PSHRQB, EXTRACTPS, PEXTRB, PEXTRD, PEXTRW, PEXTRQ	These new instructions simplify data insertion and extraction between QPR (or memory) and XMM registers.	Broadly useful for automated compiler vectorization of data processing written in high-level languages (like C and Fortran), and applications such as image processing, video processing, multimedia and gaming.
Packed Format Conversion	PMOVSQBW, PMOVSQBW, PMOVSQBW, PMOVSQBW, PMOVSQBW, PMOVSQBW, PMOVSQBW, PMOVSQBW, PMOVSQBW, PMOVSQBW, PMOVSQBW, PMOVSQBW	Converts from a packed integer (from XMM register or memory) to a zero- or sign-extended integer with wider type.	Broadly useful for automated compiler vectorization of data processing written in high-level languages (like C and Fortran), and applications such as image processing, video processing, multimedia and gaming.
Packed Test and Set	PTSET	Enables branching from SIMD decisions to support conditionally vectorized code.	Useful for improved automated compiler vectorization of data processing, image and video processing, 3-D content creation, multimedia, and gaming.
Packed Compare for Equal	PCMPEQB, PCMPEQD	Performs SIMD compare for equality of the packed QWORDS in the destination and the source operand.	Broadly useful for automated compiler vectorization of data processing written in high-level languages (like C and Fortran), and applications such as image processing, video processing, multimedia and gaming.
Pack QWORD to Unsigned WORD	PMOVSQBW	Converts packed signed QWORDS into packed unsigned WORDs using unsigned saturation to handle overflow condition. This new instruction completes the set of other instructions in this type.	Broadly useful for automated compiler vectorization of data processing written in high-level languages (like C and Fortran), and applications such as image processing, video processing, multimedia and gaming.

Penryn also implements a new divider that impacts both integer and floating point divides using a radix-16 algorithm. The algorithm computes more bits of the result of a divide each pass (four bits per iteration vs. two bits in Conroe), decreasing divide latency.



The faster divider is a very specific enhancement that should manifest itself as a performance boost in 3D and imaging applications.




Penryn's Super Shuffle Engine should also improve SSE2, SSE3 and SSE4 applications that use a lot of shuffle operations. Cache performance is also improved slightly for misaligned stores, which should improve performance, once again, in 3D and imaging applications. Finally, there are some power enhancements made to Penryn, but these are mobile-specific and thus don't apply to any of the desktop variants.


Page 2

What do we have here today? Yorkfield

The problem with Intel's codenames these days is that we've got processor family codenames and then actual chip codenames. Penryn refers to the entire family of 45nm Core architecture products that have been/are being announced, but the actual chips themselves all have their own codenames. For example, the 45nm Penryn based quad-core Xeon processor is codenamed Harpertown. Penryn on the desktop carries two names: Yorkfield and Wolfdale.

Yorkfield is quad-core Penryn for the desktop, Wolfdale is simply dual-core. Yorkfield isn't actually a different die, because a Yorkfield chip is just made up of two Wolfdale die on the same package (just like current quad-core Kentsfield Intel CPUs). This won't change until Nehalem.





The chip that Intel is launching today is the first Yorkfield: the Core 2 Extreme QX9650. The quad-core QX9650 runs at 3.0GHz with a 1333MHz FSB, much like its predecessor the QX6850. Like all Yorkfield CPUs, the QX9650 is made up of two independent dual-core die on a single package, each one with a shared 6MB L2 cache for a total of 12MB of on-chip L2 cache.



QX9650 (left), QX6850 (right)

The QX9650 will work in a number of presently shipping motherboards; we tested ours in an ASUS P35 board, but you'll have to check with your board vendor - or check [our Penryn Compatibility article](#) - to make sure that there's BIOS/board support for it. The chip is still physically an LGA-775 processor so it'll fit in any LGA-775 socket; it's just up to the motherboard guys to implement hardware and BIOS level support for the processor.

Pricing hasn't been announced yet but we expect the QX9650 to come in at the \$999 mark, the same as previous Core 2 Extreme parts.

Test Configuration

CPU:	Intel Core 2 Extreme QX9650 (3.00GHz/1333MHz) Intel Core 2 Extreme QX6850 (3.00GHz/1333MHz)
Motherboard:	Gigabyte GA-X38-DQ6 (Intel X38)
Chipset:	Intel X38
Chipset Drivers:	Intel 8.1.1.1010 (Intel)
Hard Disk:	Seagate 7200.9 300GB SATA
Memory:	Corsair XMS2 DDR2-800 4-4-4-12 (1GB x 2)
Video Card:	NVIDIA GeForce 8800 GTX
Video Drivers:	NVIDIA ForceWare 163.75
Desktop Resolution:	1600 x 1200
OS:	Windows Vista Ultimate 32-bit

Recap: When's a Penryn?

In our first Penryn preview we laid out the launch schedule for Intel's new chips; now we can finally bring you an update with more specifics including model numbers and clock speeds. Let's first look at a table that will round out the rest of 2007:

CPU	Clock Speed	FSB	L2 Cache	Availability	Pricing
Intel Core 2 Extreme QX9650	3.00GHz	1333	6MBx2	Nov 12	\$999
Intel Core 2 Extreme	3.00GHz	1333	4MBx2	Now	\$999

QX6850					
Intel Core 2 Extreme QX6800	2.93GHz	1066	4MBx2	Now	\$999
Intel Core 2 Quad Q6700	2.66GHz	1066	4MBx2	Now	\$530
Intel Core 2 Quad Q6600	2.40GHz	1066	4MBx2	Now	\$266
Intel Core 2 Duo E6850	3.00GHz	1333	4MB	Now	\$266
Intel Core 2 Duo E6750	2.66GHz	1333	4MB	Now	\$183
Intel Core 2 Duo E6550	2.33GHz	1333	4MB	Now	\$163
Intel Core 2 Duo E6540	2.33GHz	1333	4MB	Now	\$163
Intel Core 2 Duo E4600	2.40GHz	800	2MB	Q4	\$133
Intel Core 2 Duo E4500	2.20GHz	800	2MB	Now	\$133
Intel Core 2 Duo E4400	2.00GHz	800	2MB	Now	\$113
Intel Pentium E2180	2.00GHz	800	1MB	Now	\$84
Intel Pentium E2160	1.80GHz	800	1MB	Now	\$84
Intel Pentium E2140	1.60GHz	800	1MB	Now	\$74

The big new introduction here is the Core 2 Extreme QX9650, the very first Yorkfield and the first Penryn we'll see on the desktop. The QX9650 will officially launch on November 12 and although Intel hasn't revealed pricing, we're expecting it to be at \$999 thus replacing the QX6850. Given that it costs Intel less money to make than a QX6850, we'd expect Intel would want to sell more of the QX9650 anyways, and pricing it more than \$999 just isn't going to help that cause.

The more interesting table however is what happens starting next year, because that's where we get some of the more mainstream Penryn chips in the market:

CPU	Clock Speed	FSB	L2 Cache	Availability	Replaces?
Bloomfield	TBD	N/A	TBD	Q4 '08	TBD
Intel Core 2 Extreme QX9770	3.20GHz	1600	6MBx2	Q1 '08	QX9650
Intel Core 2 Extreme QX9650	3.00GHz	1333	6MBx2	Nov 12	\$999
Intel Core 2 Quad Q9550	2.83GHz	1333	6MBx2	Q1 '08	Q6700
Intel Core 2 Quad Q9450	2.66GHz	1333	6MBx2	Q1 '08	Q6600
Intel Core 2 Quad Q9300	2.50GHz	1333	3MBx2	Q1 '08	Q6600
Intel Core 2 Duo E8500	3.16GHz	1333	6MB	Q1 '08	E6850
Intel Core 2 Duo E8400	3.00GHz	1333	6MB	Q1 '08	E6750
Intel Core 2 Duo E8300	2.83GHz	1333	6MB	Q2 '08	E8200
Intel Core 2 Duo E8200	2.66GHz	1333	6MB	Q1 '08	E6550
Wolfdale 3M	TBD	1066	3MB	Q2 '08	E4700
Intel Core 2 Duo E4700	2.6GHz	800	2MB	Q1 '08	\$266

Look at those availability dates: Penryn is coming for your children in Q1 '08; only the E8300 and Wolfdale 3M parts won't be out until Q2. Clocks move up a little, the QX9770 goes to 3.20GHz thanks to a 1600MHz FSB (which will be enabled by Intel's upcoming X48 chipset - yep another one), and the E8500 brings mainstream chips up to 3.16GHz.

Now look at the "Replaces?" column to get an idea for where these things will be priced. Intel's own roadmap shows the Q9550 slotting in next to the current Q6700, which means that we may be able to find it priced at around \$530.

More interesting is that Intel seems to have segmented the affordable quad-core market a bit, by replacing the ever-popular Q6600 with two chips: a Q9450 and Q9300. While functionally identical, the Q9300 is built off of two Wolfdale 3M cores (meaning it only has 6MB total L2 cache) while the Q9450 is built off of two Wolfdale 6M cores giving it 12MB of total L2 cache. Obviously the 9300 will be cheaper to make, so we'd expect to see that at or below the \$266 price point of the Q6600. The Q9450 would slot in right above the 9300 in pricing.


We would hope that Intel will price the Q9300 below the \$266 price point (*can we have a sub-\$200 quad-core, please?*), because it will actually have less cache than the current Q6600 making it cheaper for Intel to make, but possibly reducing performance over current chips. Granted it will have all the Penryn enhancements which, as you will soon see, do improve performance but we generally like having our cake and eating it too.

The dual-core market also gets interesting, with the E8000 line replacing the current E6000 series. If Intel's pricing structure remains the same then it looks like at today's prices you'll end up with an extra 166MHz, 50% more cache, SSE4 and some other tweaks for the same money. We also have to mention how well the model numbers work out with the Core 2 products; everything is in nice increments of 100, just like when Conroe first launched. Ah those were the days....


The other important item to note on the roadmap going forward is that top line in the table - yep, the one that says Bloomfield. Bloomfield is none other than Nehalem, the 45nm successor to Penryn. It's a brand new architecture complete with an on-die memory controller, SMT (Symmetric Multi-Threading - 2 threads per core) and 8MB of shared cache (probably L3 shared among all four cores). While it's still a year away, it's very nice to see it on an Intel roadmap this far in advance of its launch.

Everything You Need to Know: Yorkfield vs. Kentsfield

The big question is: how does Yorkfield stack up to Kentsfield, the current 65nm quad-core part from Intel? Thankfully, the QX9650 runs at the same clock speed/FSB as the Kentsfield-based QX6850 making our job a little easier. We put these two bad boys up against one another and waited for the smoke to clear.




As we saw in our initial Penryn preview, the general application suite tests just don't show any real improvements of Penryn over Conroe. SYSMark 2007's overall performance went up by 2.6% but that's far from significant. If you look at the individual tests however, you'll see that the 3D benchmarks went up by 4.5% and the productivity suite went up by a similar 4.6% margin. The improvements in the 3D suite are expected, given the Radix-16 and Super Shuffle Engine enhancements to the Penryn core. The productivity suite most likely benefits from the Radix-16 divider as well as the larger cache. Overall the performance boost just isn't that significant, but given that Penryn is expected to arrive without a price premium, any performance improvement is better than nothing.



We see the same situation under PC WorldBench 6: the overall performance boost of Penryn over Conroe is basically nil. Digging deeper we see performance increases on the order of 4 - 7% if we look at the individual 3D rendering or media encoding tests.


Next up we've got our encoding tests, and given what we've already seen, we should expect some nice gains here:



And nice gains we do get. The clock-for-clock improvement under DivX 6.6 is a nice and round 10%, although curiously enough we get a meager 3% from our Windows Media Encoder test. Our QuickTime H.264 test shows a more average 4% performance increase, similar to our WME test.


Professional 3D rendering and image manipulation are also areas where we expect to see reasonably high gains from Penryn. Our 3dsmax, Cinebench and Lightwave tests remain unchanged from previous articles. The CS3 benchmark is the same Retouch Artists test we've used in previous CPU reviews. In order to

avoid having a graph scale that ranged from zero to the thousands we simply included percentage improvements here rather than actual numbers (e.g. 3dsmax scores are single digits while Cinebench scores are 4 digits):



As expected, we're seeing reasonable gains here. 3dsmax shows a 5% increase, Cinebench 8%, Lightwave 6% and Photoshop CS3 has a healthy 10% performance boost in store for us. Without a doubt encoding and 3D/image manipulation are the real strengths of Intel's Penryn architecture.

The final Penryn test is all about 3D gaming and thankfully we have a number of new titles to play with. All of our benchmarks were run at 1024 x 768 to avoid being GPU bound and with high quality settings (the exceptions being Crysis and World in Conflict which both used Medium quality defaults).



* Denotes performance in minutes, lower is better

Penryn's gaming performance is really all over the place. Titles like Oblivion and Bioshock see absolutely no performance increase, while Crysis gives us a mild 3.7% performance boost. Half-Life 2: Episode Two and Unreal Tournament 3 enjoy a 5.1% and 6.4% increase respectively, while World in Conflict and Quake Wars are both at around 9%.

If anything the gamut of gaming benchmarks sums up Penryn's performance improvement over Conroe quite well: it really varies from nothing to something.


Diving Deeper: SSE4 Performance

One of Penryn's real strengths is in its support for SSE4, which has the ability to really provide a tremendous performance advantage for some time to come. Unfortunately, as is usually the case with new instructions, it's going to take a while for applications to actually utilize them. Such is the case with SSE4 as the only benchmarks we're able to bring you come directly from Intel, but thankfully they are of real world usage models. Both tests we've actually showed you in the past, during Intel's own sanctioned Penryn previews, and both involve some sort of encoding.

The most important test is a DivX encode using VirtualDub 1.7.6 and DivX 6.7. SSE4 comes in if you choose to enable a new full search algorithm for motion estimation, which is accelerated by two SSE4 instructions: MPSADBW and PHMINPOSUW. The idea is that motion estimation (figuring out what will happen in subsequent frames of video) requires a lot of computation of sums of absolute differences, as well as finding the minimum values of the results of those computations. The SSE2 instruction PSADBW can compute two sums of differences from a pair of 16B unsigned integers; the SSE4 instruction MPSADBW can do eight.

According to Intel's own research on motion estimation with SSE4, the same search algorithm can take 71 cycles per 16x16 pixel block using the SSE2 SAD (sum of abs differences) instruction, compared to only 26 cycles using the SSE4 version. The latency reduction results in an obvious performance increase.






We used VirtualDub 1.7.6 and DivX 6.7 with SSE4 Full Search enabled to measure the impact of this motion estimation optimization. Note that the motion estimation that's taking place here is more accurate than the default DivX setting, so both SSE4 and SSE2 versions of the algorithm result in slower performance (but better quality) than with it disabled.

	SSE2 Search	SSE4 Search
Intel Core 2 Extreme QX9650 (3.0GHz)	21.9 seconds	15.1 seconds
Intel Core 2 Extreme QX6850 (3.0GHz)	35.2 seconds	N/A

On our QX9650, the full search with SSE4 enabled runs about 45% faster than with SSE2 only - impressive! Note also that the Penryn QX9650 offers better SSE2 performance in this test as well, coming in about 61% faster than the QX6850. The total performance increase from QX6850 SSE2 to QX9650 SSE4 in this test is an incredible 133%. Obviously, this is not going to be the norm in many other applications, but there's definitely some potential for meaningful optimizations in certain applications.

It's important to note that the PHMINPOSUW instructions doesn't appear to be in AMD's proposed SSE5 specification, although MPSADBW looks like it'll make it. AMD will eventually add full SSE4 support to its processors but not until the 2009/2010 time frame from what we've heard.



Our second benchmark from Intel is an MPEG-2 encode of an HD video using TMPGEnc 4.0.

	TMPEGEnc 4.0
Intel Core 2 Extreme QX9650 (3.0GHz)	103 seconds
Intel Core 2 Extreme QX6850 (3.0GHz)	135 seconds


The performance difference is a little less significant here, with the SSE4-less QX6850 taking about 31% more time to encode the input file than the QX9650.

Both of these are very real-world implementations of SSE4; unfortunately, it's tough to say how long it will be before we see widespread use of the new instructions.

I've Got the Power: 45nm vs. 65nm

Since we're dealing with the same clock speeds as Intel's 65nm processors, power consumption has definitely gone down with the move to Penryn. Let's look at this thing at idle and under load running our WME9 test:






At idle, the QX9650 draws an impressive 34W less than the QX6850 - there's 45nm high-k + metal gate transistors in action for you.

Under load the power advantage is even more impressive, with a 47W delta the QX9650 under load uses only 11W more than its predecessor at idle. If you weren't dazzled by the performance improvements of Penryn, the reduction in power consumption is worth getting excited about.

Page 7

Overclock Me Baby

Naturally we wanted to see how much was left on the table with Intel's 45nm process. We can't help but think that had Phenom been out sooner, we'd be seeing a > 3GHz Penryn launch, but how easy would it be for Intel to ramp up clock speeds?



Our unlocked QX9650 had no problems hitting 333MHz x 12.0, for a final clock speed of 4.0GHz. We had to increase the stock voltage of 1.25V up to 1.40V to achieve it, but the overclock required no additional cooling beyond the standard Intel heatsink/fan. At lower voltages, 3.66GHz should be an easy target to reach.

Page 8

Another Price Drop? A Competitive Update

Although this couldn't be further from the matter at hand, AMD has quietly dropped its prices since the last time we looked at upper mainstream performance. The new/old AMD pricing structure is as follows:


CPU	Old Price	New Price
AMD Athlon 64 FX-74	\$599/pair	\$599/pair
AMD Athlon 64 X2 6000+	\$178	\$167
AMD Athlon 64 X2 5600+	\$157	\$146
AMD Athlon 64 X2 5200+	\$136	\$125
AMD Athlon 64 X2 5000+	\$125	\$115
AMD Athlon 64 X2 4800+	\$115	\$104
AMD Athlon 64 X2 4400+	\$94	\$89
AMD Athlon 64 X2 4000+	\$73	\$68
AMD X2 BF-2350	\$91	\$96

AMD X2 BE-2300	\$73	\$91
----------------	------	------

The price changes are fairly minor, but they do change the way we have to compare these processors. For example, the Athlon 64 X2 6000+ used to be priced close to the Core 2 Duo E6750, but now it's an E6550 competitor, improving AMD's competitive stance. For some reason, AMD's 45W Athlon X2 processors actually went up in price, possibly due to fluctuation in yields.

The other major change is that below the 6000+, all of AMD's chips compete with Intel's E4000/E2000 series, not the E6000 line. We'll be working on an update to our [Midrange CPU Roundup](#) to take some of these changes into account, but for this review we'll do a quick update looking at the 6000+ compared to its new price competitor, the E6550.

The chart below shows the percent difference in scores between the 6000+ and the E6550; the blue bars mean that Intel won that test and the green bars indicate that AMD won:





AMD definitely gets more competitive with its price drop, but Intel still holds on to the competitive advantage. Not only is the E6550 faster overall, it is also a cooler running processor (keeping in mind that the 6000+ is still a 90nm core) with more overclocking headroom.

The Phenom Question

The CPU wars aren't over for the year, despite there being only 64 days left in the calendar. AMD is committed to releasing Phenom in 2007, and with it will come a definite change in the balance between AMD and Intel. No one is expecting AMD to take the overall performance crown away from Intel - AMD simply won't be able to hit the clock speeds necessary - but will it be any more competitive at mainstream price points?


When AMD's Barcelona launched we attempted to [simulate desktop performance vs. its K8 architecture in a handful of applications](#); we came up with the following chart:





Using those numbers, we took our Athlon 64 X2 6000+ results and scaled them according to where we expect Phenom to perform. We then put this simulated Phenom head-to-head with an identically clocked Core 2 Duo E6850 as well as the price-competitor to the 6000+, the E6550.

Note that this is a very rough comparison, first because the scaling values we have were taken from a quad-core Barcelona vs. two dual-core K8s and we're applying it to a dual-core K8 here. Second, AMD isn't going to be launching a 3.0GHz dual-core Phenom until sometime next year - most likely after Intel has already started shipping mainstream Wolfdale parts - which means that AMD will be competing against a completely different beast when that happens. Regardless, the comparison is still an interesting one to make; let's see if we can get any sort of expectations for what is to come.



*Denotes time in seconds, lower bars mean better performance

Clock for clock, Intel still holds on to the lead, but note that there are some situations where our simulated 3.0GHz Phenom X2 outperforms the E6550, which could be its price competitor. Again, we're assuming that a dual-core desktop Phenom will scale as well as our quad-core Barcelona and we're also assuming that AMD can price Phenom this competitively and that Intel doesn't respond with even more aggressive pricing. That's a lot of assumptions for AMD to be able to pull ahead of Intel at the same price point, but then we need to factor in the other P: Penryn.

In a couple of the areas that will be close between AMD and Intel, such as DivX and 3dsmax, Penryn happens to do really well. That means any gains AMD could make there with Phenom may be negated by the performance boost we're seeing from Penryn, leaving us in much of the same situation that we are in today. The only saving grace for AMD is that there are some areas where Intel just doesn't get a big performance boost from Penryn (e.g. SYSMark, Lightwave), and in those benchmarks Phenom will make AMD more competitive.

The end result is that we expect Phenom to make AMD more competitive, but because of Penryn and aggressive pricing - and if our assumptions are correct - it doesn't look like we'll see the sort of upset that Intel pulled on AMD with the launch of the original Core 2. (Or that AMD pulled on Intel with K8 vs. NetBurst.)

Page 10


Final Words

It's almost a bit disheartening to see Penryn launch at only 3.0GHz, as we know the architecture is capable of so much more. Unfortunately, what it looks like we're seeing here today is an artificial slowing of Intel's roadmap in response to the competition. While the tick-tock model is still very much in play and Nehalem is on track for a late '08 release, we're getting a 1333MHz FSB 3.0GHz QX9650 today instead of a 1600MHz FSB 3.2GHz Q9770 because there's no competitive reason for anything more. AMD's Phenom still isn't out and all indications are that we won't see 3GHz from AMD until sometime next year, meaning that Intel definitely doesn't have any real incentive to push the performance envelope this year.

Penryn really offers a wide range of performance; at worst it's no faster than its predecessor, but at best we're looking at low double digit performance gains. The biggest improvements seem to come in the encoding, 3D rendering and image manipulation applications, which makes sense when you look at the nature of the architectural improvements in Penryn.

The tremendous reduction in power consumption is very exciting and arguably one of the more desirable features of Penryn. The initial SSE4 results do look promising, but as we've often seen with new instructions, it's anyone's guess as to when we'll see widespread adoption of them from developers (and how much they will truly help in everyday use). The fact that we're able to bring you real world DivX performance results today though is promising.

Given that we're expecting a seamless price transition, it's tough to really find fault with Penryn. While the performance bumps can vary from minor to major, the cost should be nothing to the end user, in which case we can't complain. It's not the revolutionary step that Conroe was, but Penryn is merely making everything a little better all while driving power consumption down. So long as Intel keeps its prices the same, we're happy.





So what are the final words on it? As an evolution of Conroe, Penryn gets our vote. We just want it at lower price points, please.

Updated CPU Cheatsheet - Seven Years of Covert CPU Operations

Date: Aug 28, 2004
Type: CPU & Chipset
Manufacturer: Various
Author: Jarred Walton
Page 1

Introduction

Update: 8/27/04 - The charts have all been revised. Thanks go out to all the people that posted corrections in the comments section as well as sending them via email. In addition to the corrections, some further information and commentary has been added to the pages. For anyone that actually comes back to this article for reference information, enjoy the changes!

Foreword by Kristopher Kubicki:

From time to time we stumble upon some truly gifted and patient people here at AnandTech. Some weeks ago I wrote a [CPU codename cheatsheet](#) as just something to do in an airport terminal to kill time. Very soon after, an extremely diligent Jarred Walton showed me his rendition of the CPU family tree that he was keeping just for fun! Knowing I was bested, I offered Jarred a chance at writing a pilot for AT, and here it is! Please enjoy the second, extremely thorough CPU Cheatsheet 2.0.

*But loud! what lurks in yonder chassis, hot?
A CPU, my programs it will run!
O Pentium, Pentium! wherefore art thou Pentium?
Obscure thy benchmarks and refuse thy name.
What's in a name? that which we call a chip
By any other name would run as fast.*

My sincere apologies to Shakespeare, but that mangled version of *Romeo and Juliet* is an apt description of the world of computer processors. Once upon a time, we dealt with part numbers and megahertz. Larger numbers meant you had a faster computer. 80286 was faster than 8088 and 8086, and the 80386 was faster still, with the 80486 being the king of performance. Life was simple, and life was good. But that is the distant past; welcome to the present.

Where once we had a relatively small number of processor parts to choose from, we are now inundated with product names, model numbers, code names, and features. Keeping track of what each one means is becoming a rather daunting task. Sure, you can always try Googling the information, but sometimes you'll get conflicting information, or unrelated web sites, or only small tidbits of what you're trying to find out. So, why not put together a clear, concise document that contains all of the relevant information? Easier said than done; however, that is exactly what is attempted in this article.

In order to keep things even remotely concise, the cutoff line has been arbitrarily set to the Pentium II and later Intel processors, and the Athlon and later AMD processors. Anything before that might be interesting for those looking at the history of processors, but for all practical purposes, CPUs that old are no longer worth using. Also absent will be figures for power draw and heat dissipation, mainly because I'm not overly concerned with those values, not to mention that AMD and Intel have very different ways of reporting this information. Besides, Intel and AMD design and test their CPUs with a variety of heatsinks, motherboards, and other components to ensure that everything runs properly, so if you use the proper components, you should be fine.

So what will be included? For this first installment, details on clock frequencies, bus speeds, cache sizes, transistor counts, code names, and a few other items has been compiled. The use of model numbers with processors is also something people will likely have trouble keeping straight, so the details of processors for all Athlon XP and later AMD chips and Pentium 4 and later Intel chips will follow. The code names and features will be presented first, with individual processor specifics listed on the later pages. As a whole, it should be a useful quick reference - or cheat sheet, if you prefer - for anyone trying to find details on a modern x86 processor.

With that said, on to the AMD processors. Why AMD first? Because someone has to be first, and AMD comes before Intel in the alphabet.

Page 2

AMD Cheat Sheet

AMD Processors									
Argon (K7)	Athlon	Slot A	500-700	512K	22 + cache	250	184	100	
Pluto (K75)	Athlon	Slot A	550-850	512K	22 + cache	180	102	100	
Orion (K75)	Athlon	Slot A	900-1000	512K	22 + cache	180	102	100	
Spitfire	Duron	462	800-950	64K	25	180	100	100	
Morgan	Duron	462	900-1300	64K	25.2	180	106	100	
Thunderbird	Athlon "B"	462	650-1400	256K	37	180	117	100	
Thunderbird	Athlon "C"	462	1000-1400	256K	37	180	117	133	
Palomino	Athlon XP/M	462	850-1733	256K	37.5	180	129	100/133	
Palomino	Athlon MP	462	1000-1733	256K	37.5	180	129	100/133	1-2
Thoroughbred A	Athlon XP	462	1467-1833?	256K	37.5	130	80	133	
Thoroughbred B	Athlon XP/M	462	1200-2133	256K	37.5	130	84	133	
Thoroughbred B	Athlon XP	462	2083-2250	256K	37.5	130	84	166	
Thoroughbred B	Athlon MP	462	1667-2133	256K	37.5	130	84	133	1-2
Barton	Athlon XP/M	462	1467-2133	512K	54.3	130	101	133	
Barton	Athlon XP	462	1833-2167	512K	54.3	130	101	166	
Barton	Athlon XP	462	2100-2200	512K	54.3	130	101	200	
Barton	Athlon MP	462	2133	512K	54.3	130	101	166	1-2
Applebred	Duron	462	1400-1800	64K	25.2*	130	84*	133	
Thorton	Athlon XP	462	1667-2067	256K	37.5*	130	101*	133	
Thoroughbred B	Sempron	462	1500-2000+	256K	37.5	130	84	166	
Sledgehammer	Athlon FX	940	2200-???	1024K	105.9	130 SOI	193	200	Y
Sledgehammer	Opteron	940	1400-2400	1024K	105.9	130 SOI	193	200	Y 1-8
Sledgehammer	Athlon FX	939	2400-???	1024K	105.9	130 SOI	193	200	Y
Clawhammer	Athlon 64	754	1800-2200(?)	512K	105.9	130 SOI	193	200	Y
Clawhammer	Athlon 64	754	2000-2400(?)	1024K	105.9	130 SOI	193	200	Y
Newcastle	Athlon 64	754	1800-2600(?)	512K	68.5	130 SOI	144	200	Y
Newcastle	Athlon 64	939	2200-2600(?)	512K	68.5	130 SOI	144	200	Y
San Diego	Athlon FX	939	2600-???	1024K	105.9(?)	90 SOI	114(?)	200	Y
Paris	Sempron	754	1800-???	256K	-50(?)	130 SOI	118	200	N
Venus	Opteron 1xx	940				90 SOI	200?	Y	
Troy	Opteron 2xx	940				90 SOI	200?	Y 1-2	
Athens	Opteron 8xx	940				90 SOI	200?	Y 1-8	
Odessa	Athlon 64 M?	754?		512K		130 SOI	200?	Y	
Winchester	Athlon 64	939		512K	68.5(?)	90 SOI	83(?)	200	Y
Dublin	Athlon XP-M	462		37.5		130 SOI	128	200?	N
Newark	Athlon 64-M LP	754?				90 SOI	200?	Y	
Lancaster	Athlon 64 M	754?				90 SOI	200?	Y	
Georgetown	Athlon XP M	462/754?				90 SOI	200?	N?	
Sonora	Athlon XP-M LP	462/754?				90 SOI	200?	N?	
Denmark	Opteron 1xx	940				90 SOI	200?	Y	
Italy	Opteron 2xx	940				90 SOI	200?	Y 1-2	
Egypt	Opteron 8xx	940				90 SOI	200?	Y 1-8	
Toledo	Dual Core FX	939				90 SOI	200?	Y 2C	
Palermo	Sempron (?)	939 (?)		256K?	-50(?)	90 SOI	62(?)	200	N?
Oakville	Athlon 64 Mobile	754?		512K?		90 SOI	200?	Y	

Victoria	Sempron (?)	754 (?)	256K? ~50(?)	90 SOI	62(?)	200	N?
----------	-------------	---------	--------------	--------	-------	-----	----

* Die Size and/or transistor count is based off a larger CPU core with a portion of the die disabled.

** Various steppings/sources listed different die sizes.

*** The bus speed all Athons/Durons is double-pumped, but the CPU multiplier is based off the listed speed.

A few notes to clarify the information. The stated die sizes and transistor counts for the Applebred and Thorton reflect the fact that these processors are Thoroughbred and Barton cores, respectively, with half of the L2 cache disabled, which is why they have a single asterisk next to them. There have been reports of hacking the Thorton processors and turning them into full Barton CPUs, but considering the insignificant cost difference these days, it's probably not worth worrying about. AMD plans on discontinuing the Barton soon anyway, and will use the old Thoroughbred core for the Socket A Sempron chips.

Transistor counts on Paris, Victoria, and Palermo are likely off, but it remains to be seen how AMD actually configures these chips. Early Athlon 64 512K cache chips for socket 754 were Clawhammer cores with half the cache disabled, but the newer models (i.e. 3200+ at 2.2 GHz with 512K, 3400+ 2.4 GHz 512K, and 3700+ 2.6 GHz with 512K) appear to be actual Newcastle cores. The same could very well happen with the Paris cores, where initial shipments are "downgraded" Newcastle cores, and later versions may physically remove the ~18.7 million transistors used in the L2 cache. Regardless, values on these cores should be taken with a grain of salt.

Unreleased processors will likely change from these current estimates, and question marks indicate best guess data at present. If you notice any errors or if you have additional information on forthcoming processors, let us know in the comments section or email.

Take note of the Toledo, Denmark, Italy, and Egypt cores; the 2C next to it stands for dual core. All four models use the same basic core and should come out around the same time in early 2005. Whether they launch as planned remains to be seen, and precise details about the internal layout are not yet clear - recent news suggests that each core will have its own L2 cache. Dual core is best described as SMP on a single chip, and while on the subject of SMP, please note that all of the Athlon XP processors could support multi-processor configurations *unofficially*. 2-way SMP was almost a certainty, but none of the CPUs were verified to function in such a configuration by AMD. While it would not be prudent to take such a risk as a business, quite a few enthusiasts saved themselves a lot of money by putting XP chips into SMP motherboards instead of spending the extra money on MP chips.

The basic core of the Athlon, from the Pluto all the way through the latest Newcastle and Paris processors, changed very little since its inception. It has a 10 stage integer pipeline and 15 stage floating point pipeline, with three identical Arithmetic/Logic Units (ALUs), Address Generation Units (AGUs), and Floating Point Units (FPUs). The FPUs also handle the MMX, 3DNow!+, and SSE/SSE2 support. Opteron increased the length to 12/17 stages, in addition to bringing 64-bit support. Future versions of the Athlon 64 will likely increase the length of the pipeline past the current 12/17 stages in order to increase clock speeds, but I doubt that AMD will ever show the *hubris* of Intel by creating a 31 stage pipeline - at least, not on any iteration of the Athlon architecture. This is especially a problem with the increasing power leakage of high clockspeeds and increasingly small process technology. Until those issues are resolved, I think it's safe to say that pipeline lengths will stay in the 10 to 15 stages (for integers) range with AMD.

Update: One reader was good enough to send a link to AMD's site where they actually list the Opteron as being a 12/17 design. (Thanks Tom!) Finding any good details on the Intel and AMD sites can be a major chore, most likely due to the level of competition between the companies as well as their size. There's a rule somewhere that the larger a company gets, the less informative and helpful their web site becomes! For those that want the link, here's the [Opteron information](#). That means that all Athlon 64 designs are also 12/17, of course. The Denmark, Italy, and Egypt CPUs are also dual core, it appears, and their entries have been updated to reflect this. (The old roadmap didn't include that information.)

Intel Cheat Sheet

				L1 Cache	L3 Cache	Bus Speed (MHz)	
Intel IA32/EM64T Processors							
Covington	Cel	Slot 1	266/300	8K+8K	7.5	350 118	66
Mendocino	Cel ("A")	Slot 1	266-433	16K+16K 128K	19	250 154	66
Mendocino	Cel ("A")	370	233-533	16K+16K 128K	19	250 154	66
Coppermine-128	Cel ("A")	370	533-766	16K+16K 128K	28*	180 106	66
Coppermine-128	Cel ("A")	370	800-1100	16K+16K 128K	28*	180 106	100
Klamath	P II	Slot 1	233-333	16K+16K 512K	7.5+37.2	350 203+L2	66
Deschutes	P II	Slot 1	266-333	16K+16K 512K	7.5+37.2	250 118+L2	66
Deschutes	P II	Slot 1	350-450	16K+16K 512K	7.5+37.2	250 118+L2	100
Deschutes	P II Xeon	Slot 2	400-450	16K+16K 512K	7.5+37.2	250 118+L2	100
Deschutes	P II Xeon	Slot 2	400-450	16K+16K 1M	7.5+74.4	250 118+L2	100
Deschutes	P II Xeon	Slot 2	450	16K+16K 2M	7.5+148.8	250 118+L2	100
Katmai	P III	Slot 1	450-600	16K+16K 512K	9.5+37.2	250 131+L2	100
Katmai	P III B	Slot 1	533-600	16K+16K 512K	9.5+37.2	250 131+L2	133
Tanner	P III Xeon	Slot 2	500, 550	16K+16K 512K	9.5+37.2	250 128+L2	100
Tanner	P III Xeon	Slot 2	500, 550	16K+16K 1M	9.5+74.4	250 128+L2	100
Tanner	P III Xeon	Slot 2	500, 550	16K+16K 2M	9.5+148.8	250 128+L2	100
Cascades**	P III Xeon	Slot 2	600-1000	16K+16K 256K	28.1	180 106-90	133
Cascades	P III Xeon	Slot 2	700	16K+16K 1M		180 210?	100
Cascades	P III Xeon	Slot 2	700, 900	16K+16K 2M		180 385	100
Coppermine**	P III	Slot 1	550-1000	16K+16K 256K	28.1	180 106-90	100
Coppermine**	P III B	Slot 1	533-1000	16K+16K 256K	28.1	180 106-90	133
Coppermine**	P III E	370	500-1100	16K+16K 256K	28.1	180 106-90	100
Coppermine**	P III EB	370	533-1133	16K+16K 256K	28.1	180 106-90	133
Tualatin	Cel ("A")	370	1000-1400	16K+16K 256K	28.1	130 80	100
Tualatin	P III	370	1000-1333	16K+16K 256K	28.1	130 80	133
Tualatin	P III S	370	1133-1400	16K+16K 512K	45.9	130 110?	133
Willamette	Cel-128	478	1700-1800	12Ku+8K 128K	36.5	180 217*	100
Willamette	P 4	423	1300-2000	12Ku+8K 256K	42	180 217	100
Willamette	P 4	478	1500-2400	12Ku+8K 256K	42	180 217	100
Foster	Xeon DP	603	1400-2000	12Ku+8K 256K	42	180 217	100
Foster	Xeon MP	603	1400, 1500	12Ku+8K 256K 512K 42+37?	180		100
Foster	Xeon MP	603	1600	12Ku+8K 256K 1M 42+74?	180		100
Northwood	Cel	478	1400-2800	12Ku+8K 128K	36.5	130 131?	100
Northwood	Mob. Cel.	478	1400-2800	12Ku+8K 256K		130	100
Northwood**	P 4	478	1800-2600	12Ku+8K 512K	55	130 146-131	100
Northwood**	P 4 "B"	478	2267-2800	12Ku+8K 512K	55	130 146-131	133
Northwood**	P 4 HTT	478	3067	12Ku+8K 512K	55	130 146-131	133
Northwood**	P 4 "C"	478	2400-3400	12Ku+8K 512K	55	130 146-131	200
Gallatin**	P 4 EE	478	3200-3400	12Ku+8K 512K 2M	55+123	130 231-237?	200
Prestonia	Xeon DP	603	1600-3000	12Ku+8K 512K	55	130	100
Prestonia	Xeon DP	604	2000-3067	12Ku+8K 512K	55	130	133
Prestonia	Xeon DP	604	3067-3200	12Ku+8K 512K 1M	55+61	130	133
Gallatin	Xeon MP	603	1500-2800	12Ku+8K 512K 1M	55+61	130	100
Gallatin**	Xeon MP	603	2000-2700	12Ku+8K 512K 2M	55+123	130 231-237?	100
Gallatin	Xeon MP	603	3000	12Ku+8K 512K 4M	55+246?	130	100
Prescott 256?	Cel D	478/775	2400-3200	12Ku+16K 256K		90	133
Prescott	P 4 "A"	478	2400-2800	12Ku+16K 1M	125	90 112	133
Prescott	P 4 "E"	478	2800-3400	12Ku+16K 1M	125	90 112	200
Prescott	P 4 "E"	T/775	2800-???	12Ku+16K 1M	125	90 112	200
Prescott	P 4 "E"	T/775	??-???	12Ku+16K 2M		90	200/266
Nocona	Xeon	T/775?	2800-3600+	12Ku+16K 1M	125	90 112?	200
Irindale					2M	90	200?
Banias	Cel M	478M	1300-1500	32K+32K 512K		130	100
Banias	P M	478M	900-1800	32K+32K 1M		130	100
Dothan	Cel M	478M	900-1500	32K+32K 1M		90	100/133
Dothan	P M	478M	1000-2400	32K+32K 2M		90	100/133
Potomac						65	
Smithfield							2C
Jonah	P M?					65?	2C

Tulsa
Merom
Conroe
Gilo
Whitefield

			L1 Cache	L3 Cache				Bus Speed (MHz)	SMP***
Intel IA64 Processors									
Merced****	Itanium1	PAC-418	733-800	16K+16K	96K	2-4M	25+300	180 300	66 512
McKinley+	Itanium2	PAC-611	900-1000	16K+16K	256K	1.5-3M	221	180 421	100 512
Deerfield	Itanium2	PAC-611	1000-1500?	16K+16K	256K	1.5M?		130 266?	100 512
Madison++	Itanium2	PAC-611	1300-1500?	16K+16K	256K	2-6M	477	130 374	100 512
Fanwood	Itanium2	PAC-611	1500-1667?	16K+16K	256K	9M		130 100/166	512
Montecito	Itanium2?					24M?	1700?	90	2C?
Millington	Itanium2?								
Dimona	Itanium2?								2C
Montvale	Itanium2?								
Tukwila	Itanium2?								16C?
Foxton	Itanium2?								
Pellston	Itanium2?								

* Die Size and/or transistor count is based off a larger CPU core with a portion of the die disabled.

** Various steppings/sources listed different die sizes.

*** The bus speed on the P4, PM, CM, and Itanium is quad-pumped, but the CPU multiplier is based off the listed speed.

**** Figures for Merced based off of 4M L3 cache version.

+ Figures for McKinley based off 3M L3 cache version.

++ Figures for Madison based off 6M L3 cache version.

+++ All Itaniums are said to be 512-way SMP capable, but this is more a factor of the motherboard and system design than the chip itself (I think).

Notes on the Intel side of things are similar to the AMD side. There are again a couple cores that have an asterisk, indicating that the core was a "downgraded" version of a faster core, mostly with the Celeron processors. The double-asterisks are for chips that had varying die sizes in the various steppings. This probably occurs to a small degree in most chips, but in the Cascades, Coppermine, and Northwood cores, the changes were well documented and rather drastic. Thoroughbred A to B in AMD was only a 4 mm2 die size increase, while Coppermine fluctuated between 106 mm2 to 90 mm2, and Northwood went from 146 mm2 to 131 mm2. My guess is that it was due in part to hand-optimizing the layouts of the cores, but if anyone has precise details on the hows and whys of the decreases, I would like to hear them.

In order to make the charts fit nicely within the space constraints, x86-64 was removed from the column lists. As of now, the only Intel CPUs that are known to include x86-64 support are the Nocona and Potomac cores. There will almost certainly be more in the future. The L1 cache of the P4 chips includes a trace cache, which stores decoded micro-ops, abbreviated uops. In the chart above, the trace cache corresponds to the L1 instruction cache found in typical CPUs, and 12Ku+16K means the cache has the ability to store 12,000 micro-ops as well as a standard 16KB of L1 data cache.

You can see that Intel also has 2C (dual core) designs in their roadmap, as well as a highly speculative 16C (sixteen core!) Itanium. Whether or not Tukwila will ever see the light of day is anyone's guess - it could simply be a mythical design that some hardware sites fantasize about. Transistor count on such a chips would likely be several BILLION transistors. (On a different note, I was recently up in Tukwila, WA purchasing a mountain bike from a pawn shop. They didn't have any processors for sale, unfortunately.)

In contrast to AMD, Intel has had several major architecture revisions during the past seven or so years. AMD pretty much stuck with the K7/Athlon core for all their processors, which was admittedly a very good design. Intel, with its deeper pockets, attacked on numerous fronts. First was the Pentium III line, which more or less ended in a draw with their rival AMD. Prompted by marketing - because "clockspeed sells" - Intel came up with a radical new architecture dubbed NetBurst, the basis of the Pentium 4. NetBurst was a success on the desktop, but it really was too power hungry for laptops, so Intel decided to pursue a completely separate architecture for its mobile processors, which is now also penetrating Blade and other low voltage markets. Finally, shortly after the launch of the Athlon 64, Intel countered with their reworked NetBurst architecture and the Prescott line of processors. Add to this the long-awaited launch of IA-64 (roughly ten years in the making!) which was a completely new architecture that was even more radical than NetBurst. Intel has been busy, needless to say.

For their desktop chips, SMP was available both officially and unofficially. The Celeron chips were not intended for SMP use and were never validated (by Intel) to work in such configurations. However, enterprising motherboard makers like Abit with their BP6 board allowed users to run early Celerons in dual CPU configurations. Intel put a stop to that with Coppermine-128 and Tualatin-256 (if you can call it that) Celerons. The P3 Xeon chips were all "multi-processor" configurations, capable of up to 8-way SMP. Such support was more dependent on the motherboard and chipset, though, so most setups topped out at 4-way SMP. Intel had a chipset that linked two 4-way buses together for their 8-way configuration, while ServerWorks created a chipset and motherboard that supported 8-way directly. In theory, they could have even followed Intel's example and linked two buses together to have a 16-way SMP setup, although at that point motherboard size becomes a difficult issue.

Itanium and SMP is a special case that needs further clarification. SMP is not always listed in the above chart, but all Itaniums are said to be capable of 512-way SMP. This is really more of a factor of the motherboard(s) and system design than the chip itself. For example, special high-end clustered systems have been built using AMD Athlon MP and Opteron CPUs as well as Xeon chips that have as many as 128 chips in a "single" system. Itanium is a similar case with SMP. Motherboards with up to eight sockets exist for Itanium, but 512-way SMP requires special hardware beyond the motherboard. (Please feel free to correct me if that's wrong, but I'm pretty sure this is the case. I can't imagine what a motherboard for 512 Itaniums would even look like if it were to exist - 8x8 feet in size?)

Update: A couple people pointed out issues with the naming of the Celeron processors. At the time, Intel used "A" to designate processors that overlapped an existing model. So there were cacheless Celeron 266/300 processors, and the 266/300 with 128K L2 cache had an "A" suffix. This occurred again with the Celeron 533, and once more with the Celeron 1000/1100. In a similar vein, the Klamath core was only 350 nm, while Deschutes was 250 nm. It was initially listed as 350/250 as there were certain Deschutes cores that were released as a pseudo-Klamath, for instance the P2 300 MHz SL2W8. There was not any way to actually tell (other than word of mouth) which P2 chips had the Klamath core and which had the Deschutes core. The chart has now been corrected by putting in a 250 nm 266-333 Deschutes line.

Page 4

Introduction to the Processor Charts

Before we get to the actual charts, I want to take a minute to make clear how the charts are organized. Due to the number of features involved with modern processors, it can become difficult to determine which CPU is actually faster when comparing different models. For example, do you go with the 2250 MHz Athlon XP using the Thoroughbred core, which has a 2800+ model number, or should you go with the 2000 MHz Athlon XP that uses the Barton core, which also has a 2800+ model number? With Intel, it can be even more difficult: you have different cache sizes, bus speeds, and even architectures.

Since I figure a lot of people may actually find some sort of relative sorting useful, I have attempted to do this. How you wish to rate the various factors is of course a topic that could be debated *ad nauseum*. What I am presenting is by no means a definitive answer on which model is faster, but it should give a rough estimate. Below are the various families of processors and the weighting values that I used. I then took the weight factor and multiplied that by the actual clock speed to come up with a final performance ranking.

Since this is simply a rough estimate on my part, I am not including these ranking values in the actual charts, but they are how I sorted the data. Really, the reason for their existence was to get a sorting function that more or less agreed with my own personal opinion, so if I happen to have missed a processor, or if a new processor is released, I can simply add in the processor(s) to the chart and resort it. I'm open for suggestions on how these ratings might be improved, but please realize that there will never be a definitive formula, as relative performance depends on what specific code you are running.

If you don't like math or don't really care to know precisely *how* the charts are sorted, feel free to just skip to the next page. This is only for people that really want to know details. Also, the weighting factors are within each family - they have no correlation with other processor families. (So don't get upset that the Dothan has a 1.6 weighting and Athlon FX only has 1.15!) With that said, here are the weighting factors that I used.

Duron, Athlon, Athlon XP and Sempron

64K L2 + 100 MHz bus = 0.7
64K L2 + 133 MHz bus = 0.75
256K L2 + 100 MHz bus = 0.8
256K L2 + 133 MHz bus = 0.85
256K L2 + 166 MHz bus = 0.9
512K L2 + 133 MHz bus = 0.95
512K L2 + 166 MHz bus = 1.0
512K L2 + 200 MHz bus = 1.05

Athlon 64

256K L2 + single-channel (Socket 754) = 0.9
512K L2 + single-channel (Socket 754) = 0.95
1024K L2 + single-channel (Socket 754) = 1.0
512K L2 + dual-channel (Socket 939) = 1.05

128K L2 + dual-channel (Socket 940) = 1.04
1024K L2 + dual-channel (Socket 940) = 1.11
1024K L2 + dual-channel (Socket 939) = 1.15

Celeron 2 and Pentium 4

128K L2 + 400 FSB = 0.6
256K L2 + 400 FSB = 0.75
256K L2 + 533 FSB = 0.80
512K L2 + 400 FSB = 0.84
512K L2 + 533 FSB = 0.91
1024K L2 + 533 FSB = 0.93
1024K L2 + 800 FSB = 0.98
512K L2 + 800 FSB = 1.0
512K L2 + 800 FSB + 2048K L3 = 1.15
2048K L2 + 1066 FSB = 1.2

Mobile Celeron, Mobile P4, Celeron M and Pentium M

128K L2 + 400 FSB = 0.6
256K L2 + 400 FSB = 0.75
256K L2 + 533 FSB = 0.80
512K L2 + 533 FSB + Northwood = 0.91
1024K L2 + 533 FSB + Prescott = 0.93
512K L2 + 400 FSB + Dothan = 1.25
512K L2 + 400 FSB + Banias = 1.3
1024K L2 + 400 FSB + Dothan = 1.35
1024K L2 + 400 FSB + Banias = 1.4
2048K L2 + 400 FSB = 1.5
2048K L2 + 533 FSB = 1.6

Duron and Athlon

I won't bother going into details of the early Athlon and Duron processors. They were great in their day, but they're getting to be rather long in the tooth. If there is a strong demand for more details on these processors, I will add them at a later point, but for now I simply recommend that you bite the bullet and upgrade.

For those interested in some historical information, here are a few more tidbits. The early Argon, Pluto and Orion Athlon chips had L2 cache chips contained within the Slot A cartridge. This cache could run at 1/2, 2/5, or 1/3 of the core clock speed - the faster the core, the lower the ratio. This led to situations where, for example, a 700 MHz Athlon with 350 MHz L2 would outperform the more expensive 750/300, or the 850/340 would beat the 900/300 due to the slower cache. Generally speaking, performance comparisons between the Athlon and Pentium III chips of the day were neck-and-neck affairs, with each side winning some benchmarks. Athlon had better x87 floating point performance, while Intel generally won out with features like MMX and SSE - at least in applications that were properly optimized.

The socket A processors switched to an integrated full-speed L2 cache, but the cache was half as large. The increased speed and reduced latencies, however, more than made up for the decrease in cache size. At this time, AMD was able to actually surpass Intel in raw performance for a period of time. The Athlon Thunderbird eventually reached 1.4 GHz, while the Pentium III tried for 1.13 GHz and failed. Later versions of the Pentium III dubbed Tualatin would eventually reach 1.4 GHz, but those only came after the introduction of the Pentium 4. Athlon during these times was the chip for gaming systems.

One other item worth noting is that all of the Athlon and Duron systems used the EV6 bus protocol acquired from DEC/Alpha. This was a double-pumped system bus, which improved performance relative to older buses like that used in P6 motherboards. The bus speeds listed in the charts are the base bus speed, which is then multiplied by the CPU multiplier to arrive at the final CPU speed. However, due to the double-pumping, many motherboards will list the bus speed as the doubled value. The actual performance increase gained from the doubling of the bandwidth is not as large as some might expect, but it probably accounts for somewhere between 5 to 15 percent of the total performance of the architecture, depending on the application.

The Athlon 64 and Opteron processors, meanwhile, have switched to a HyperTransport bus running at 800 MHz on the early chips and 1 GHz on socket 939 chips. The main benefit of the HT bus is that it doesn't require as many traces (wires), so it makes motherboard layouts somewhat easier to design. This also allows for multiple high-speed bus connections when used in SMP systems without resorting to designs with more layers.

Athlon XP and Sempron Processors

					Multiplier
Athlon XP (Desktop) & Sempron (Desktop Value)					
Athlon XP 1500+	1333	Palomino	256	133.3	10.0X
Athlon XP 1600+	1400	Palomino	256	133.3	10.5X
Athlon XP 1700+	1467	Palomino/TBA	256	133.3	11.0X
Athlon XP 1800+	1533	Palomino/TBA	256	133.3	11.5X
Sempron 2200+	1500	Thoroughbred B	256	166.7	9.0X
Athlon XP 1900+	1600	Palomino/TBA	256	133.3	12.0X
Athlon XP 2000+	1667	Palomino/TBA	256	133.3	12.5X
Athlon XP 2000+	1667	Thorton	256	133.3	12.5X
Athlon XP 2000+	1533	Barton	512	133.3	11.5X
Athlon XP 2100+	1733	Palomino/TBA	256	133.3	13.0X
Sempron 2400+	1667	Thoroughbred B	256	166.7	10.0X
Athlon XP 2200+	1800	TBA/TBB	256	133.3	13.5X
Athlon XP 2200+	1800	Thorton	256	133.3	13.5X
Sempron 2500+	1750	Thoroughbred B	256	166.7	10.5X
Athlon XP 2200+	1667	Barton	512	133.3	12.5X
Sempron 2600+	1833	Thoroughbred B	256	166.7	11.0X
Athlon XP 2400+	2000	Thoroughbred B	256	133.3	15.0X
Athlon XP 2400+	2000	Thorton	256	133.3	15.0X
Athlon XP 2400+	1800	Barton	512	133.3	13.5X
Athlon XP 2500+	1867	Barton	512	133.3	14.0X
Sempron 2800+	2000	Thoroughbred B	256	166.7	12.0X
Athlon XP 2600+	2133	Thoroughbred B	256	133.3	16.0X
Athlon XP 2500+	1833	Barton	512	166.7	11.0X
Athlon XP 2600+	2083	Thoroughbred B	256	166.7	12.5X
Athlon XP 2600+	2000	Barton	512	133.3	15.0X
Athlon XP 2600+	1917	Barton	512	166.7	11.5X
Athlon XP 2700+	2167	Thoroughbred B	256	166.7	13.0X
Athlon XP 2800+	2250	Thoroughbred B	256	166.7	13.5X
Athlon XP 2800+	2083	Barton	512	166.7	12.5X
Athlon XP 3000+	2167	Barton	512	166.7	13.0X
Athlon XP 3000+	2100	Barton	512	200	10.5X
Athlon XP 3200+	2200	Barton	512	200	11.0X

*** All system buses for Athlon XP, Sempron, Athlon 64, and Opteron are "double pumped", so their data rate is twice the bus speed. The multiplier is based off the listed speed.

Many of the processors listed in the charts were not commonly available, so they may not be well known. Some of these parts were shipped to OEMs who had special requirements, for example they might want to use cheaper PC2100 RAM with a Barton core. Some of the listed chips might also have been mobile parts which were mistakenly listed in the wrong table. However, the majority of these chips actually do exist in various PCs. Note also that some parts were likely to be seen more in overseas markets than in the US. If you are sure that a part is incorrect or doesn't exist, feel free to post a comment or send an email.

Athlon XP tweaked some of the finer details of the Athlon architecture to improve performance. Since XP was also going up against Pentium 4 instead of Pentium III, AMD (re)introduced model numbers and began their "clock speed isn't everything" campaign. According to AMD, the XP line was rated in terms of performance relative to the Thunderbird core, but few people actually believe that. It was almost surely market driven, as the Pentium 4 was scaling rapidly in clock speed, and the Athlon cores couldn't possibly keep up in raw MHz. And of course, AMD is correct that clock speed isn't everything - average instructions executed per clock (IPC) multiplied by clock speed would give you the real instruction throughput. Unfortunately, coming up with a precise measurement of IPC is virtually impossible - it varies depending on the code executed. Still, clock-for-clock, Athlons are definitely faster than P4 chips, and the PR ratings were relatively accurate, at least in the beginning.

As the "processor wars" continued, both companies released tweaked designs. Thoroughbred was a process shrink that brought higher clock speeds, but not as high as initially desired. A reworked Thoroughbred B core - which added an extra layer to the core, among other things - helped raise the clock limit a bit more and allowed Athlon XP to eventually reach 2250 MHz. Note that Thoroughbred B cores can often overclock to 2.3 to 2.4 GHz with sufficient cooling, while the A versions are often limited to ~2.1 GHz.

After Thoroughbred, AMD added more cache with the Barton core, and readjusted their model numbers accordingly, since more cache brought more performance. This was really where the model numbers started to become suspect, though, since Intel had also added more cache and increased bus speeds without "adjusting" any model numbers. The 2500+, 2600+ and 2800+ tended to struggle a bit in keeping up with their Intel counterparts, but the real problem came when Intel released the 200 MHz (800 FSB) "C" version of their Pentium 4. The jump to 3200+ with the 200 MHz FSB really only kept the Athlon XP competitive with the P4 2.8C in overall performance comparisons. Of course, here the model names were a stroke of genius, as many people simply assumed that a 3200+ really was the equivalent of the 3.2C.

Athlon XP-Mobile Processors

Multiplier					
Athlon XP-M (Mobile)					
Athlon XP-M 850	850	Palomino	256	100	8.5X
Athlon XP-M 900	900	Palomino	256	100	9.0X
Athlon XP-M 950	950	Palomino	256	100	9.5X
Athlon XP-M 1000	1000	Palomino	256	100	10.0X
Athlon XP-M 1100	1100	Palomino	256	100	11.0X
Athlon XP-M 1200	1200	Palomino	256	100	12.0X
Athlon XP-M 1400+	1200	Thoroughbred	256	133.3	9.0X
Athlon XP-M 1500+	1300	Palomino	256	100	13.0X
Athlon XP-M 1600+	1400	Palomino	256	100	14.0X
Athlon XP-M 1500+	1333	Thoroughbred	256	133.3	10.0X
Athlon XP-M 1600+	1400	Thoroughbred	256	133.3	10.5X
Athlon XP-M 1700+	1467	Thoroughbred	256	133.3	11.0X
Athlon XP-M 1800+	1533	Thoroughbred	256	133.3	11.5X
Athlon XP-M 1900+	1600	Thoroughbred	256	133.3	12.0X
Athlon XP-M 1900+	1467	Barton	512	133.3	11.0X
Athlon XP-M 2000+	1667	Thoroughbred	256	133.3	12.5X
Athlon XP-M 2000+	1533	Barton	512	133.3	11.5X
Athlon XP-M 2100+	1600	Barton	512	133.3	12.0X
Athlon XP-M 2200+	1800	Thoroughbred	256	133.3	13.5X
Athlon XP-M 2200+	1667	Barton	512	133.3	12.5X
Athlon XP-M 2400+	1800	Barton	512	133.3	13.5X
Athlon XP-M 2500+	1867	Barton	512	133.3	14.0X
Athlon XP-M 2600+	2000	Barton	512	133.3	15.0X
Athlon XP-M 2800+	2133	Barton	512	133.3	16.0X

*** All system buses for Athlon XP, Sempron, Athlon 64, and Opteron are "double pumped", so their data rate is twice the bus speed. The multiplier is based off the listed speed.

There's not really a whole lot to say about the Mobile AMD processors. They are identical to their desktop counterparts, except they run on lower voltages and can run at reduced clock speeds to save power. Later on, the Athlon XP-M processors gained tremendous popularity due to their unlocked multipliers, which allowed them to overclock very well, as you could keep the bus speed close to the standard 200 MHz.

There are some OEM parts as well in the Mobile Athlon market which use a different socket than the standard 462 pin socket A. For the Athlon XP, there is a 563 pin version, and for Athlon 64 there is a 638 pin version. Further details and information on these parts is, at present, lacking.

Athlon 64 and Opteron Processors

Multiplier					
Athlon 64 & "Performance" Sempron					
Sempron 3100+	1800	Paris*	256	200	9.0X 754
Athlon 64 2800+	1800	Clawhammer	512	200	9.0X 754
Athlon 64 2800+	1800	Newcastle	512	200	9.0X 754
Athlon 64 3000+	2000	Clawhammer	512	200	10.0X 754
Athlon 64 3000+	2000	Newcastle	512	200	10.0X 754
Athlon 64 3200+	2000	Clawhammer	1024	200	10.0X 754
Athlon 64 3200+	2200	Newcastle	512	200	11.0X 754
Athlon 64 3400+	2200	Clawhammer	1024	200	11.0X 754
Athlon 64 3400+	2400	Newcastle	512	200	12.0X 754
Athlon 64 3500+	2200	Newcastle	512	200	11.0X 939
Athlon 64 3700+	2400	Clawhammer	1024	200	12.0X 754
Athlon 64 FX-51	2200	Sledgehammer	1024	200	11.0X 940
Athlon 64 3700+	2600	Newcastle	512	200	13.0X 754
Athlon 64 3800+	2400	Newcastle	512	200	12.0X 939
Athlon 64 FX-53	2400	Sledgehammer	1024	200	12.0X 940
Athlon 64 FX-53	2400	Sledgehammer	1024	200	12.0X 939

Multiplier					
Opteron**					
Opteron x40	1400	Sledgehammer	1024	200	7.0X
Opteron x42	1600	Sledgehammer	1024	200	8.0X
Opteron x44	1800	Sledgehammer	1024	200	9.0X
Opteron x46	2000	Sledgehammer	1024	200	10.0X
Opteron x48	2200	Sledgehammer	1024	200	11.0X
Opteron x50	2400	Sledgehammer	1024	200	12.0X

* The Paris core does not support 64-bit computing. It is included with the Athlon 64 because of the socket and because the integrated memory controller puts it ahead of the Athlon XP in performance.

** All Opterons are available in 1xx, 2xx, and 8xx versions. x=1 is for single processor systems, x=2 is for up to dual processor systems, and x=8 is for up to octal processor systems.

*** All system buses for Athlon XP, Sempron, Athlon 64, and Opteron are "double pumped", so their data rate is twice the bus speed. The multiplier is based off the listed speed.

With the Athlon 64, as the name suggests AMD added support for 64-bit addresses and integers. This was done by widening their pathways and registers, but it wasn't a radical redesign of the core Athlon architecture. It has a pipeline that was increased to 12/17 stages, it got SSE2 support added, and the system bus was switched to a HyperTransport bus. The longer pipelines allow it to scale to somewhat higher clock speeds, and the HyperTransport buses - there are three in the Opteron - allow for better SMP, but the core remains essentially the same. The addition of x86-64 support has garnered a lot of attention, but so far it's pretty much marketing hype. It has *potential* to improve performance once 64-bit support arrives, but that *potential* has not yet been realized in the mainstream market. The scientific and academic community, however, has greeted the introduction of affordable 64-bit processing with open arms. Most consumers, meanwhile, are stuck waiting for Windows XP-64.

The reason for the superior performance of the Athlon 64 - in current 32-bit code as well as in 64-bit code to a lesser extent - lies mostly with the integrated memory controller, which *dramatically* reduces memory latencies. In effect, it helps to turn system RAM into a very large but still relatively slow L3 cache. It also continues to reduce memory latencies as clock speeds increase. Memory latencies on the Athlon XP were roughly 81 ns at 3200+, and the P4 3.2C was around 77 ns latency. Meanwhile, the Athlon 64 3400+ comes in at an astonishingly low 48 ns. As mentioned before, those latency figures are getting somewhat close to L3 cache values - for example, the L3 cache in a 3.06 GHz Xeon is about 10 ns. It's still four times slower, but it's also twice as fast as RAM on a P4 system.

No better example of this can be found than the newly introduced Paris core, a.k.a. Sempron 3100+. At 1.8 GHz, it is substantially slower than the fastest Athlon XP in core speed, and yet in typical use it outperforms even the Athlon XP 3200+. This from a part that has half as much cache as the Barton and Newcastle cores! The only area where it fails to keep up is in tasks that generally fit within the L1/L2 cache of the CPU, i.e. certain encoding tasks. In that case, the lack of raw clockspeed is a hindrance.

Of course, reduced latency isn't the entire story of the Athlon 64. In 64-bit mode, the number of useable registers for both integer and floating point operations has been doubled. Depending on the code being run, this could potentially bring 10 to 20 percent more performance. Certain applications that make heavy use of 64-bit integers can also benefit from the added 64-bit support - for example `cryptonanhv` and encoding tools.

However, MMX and SSE have provided alternative means of improving 64-bit integer performance for many years now - they just require more programming effort to realize.

Celeron, Pentium II and III Processors

I'm going to forego listing the various models of these processors for the time being. If anyone has a real desire to see them listed, feel free to let me know. If you're running one of these processors still, I can feel for you. I still use one at work, and I only upgraded from my P3 back in March. Given the price of upgrading, though - \$225 will get you a decent motherboard, 512 MB RAM, and an Athlon XP 2500+ - you really should upgrade if at all possible.

The old Pentium Pro P6 architecture was a 12-stage pipeline, more or less concluding with the Pentium III (more on that later). It had three specialized AGUs, two ALUs - one that handled simple instructions and a second one for the more complex instructions - and one FPU. The FPU also added support for SSE (which AMD lacked until the Athlon XP, but by then Intel was pushing the P4) and MMX - and they were generally faster on these instructions than AMD. That's not too surprising, considering that they created the technologies and AMD had to license them from Intel.

Intel could have certainly stuck with the design for a lot longer, as the last gasp Tualatin core offered pretty competitive performance clock for clock with the Athlon up to 1.4 GHz (the last Pentium III-S). In fact, the later 1.0A to 1.4A Celeron processors were very good overclocking chips, and a 1.1A running on a 133 MHz bus gave pretty decent performance. (I have just such a system powering my Home Theater PC.) Newer and better chipsets could have improved speed further, but Intel cut off the line and focused on pushing the Pentium 4 and NetBurst. This appears now to have been more of a marketing driven decision, although for the most part it can't be said that it was the worst idea ever.

Celeron 2 and Pentium 4 Processors

					Multiplier	L3 Cache
Pentium 4 and Celeron (Desktop)						
P4 1.3	1300	Willamette	256	100	13.0X	423
C 1.7	1700	Willamette	128	100	17.0X	478
P4 1.4	1400	Willamette	256	100	14.0X	423
P4 1.4	1400	Willamette	256	100	14.0X	478
C 1.8	1800	Willamette	128	100	18.0X	478
P4 1.5	1500	Willamette	256	100	15.0X	423
P4 1.5	1500	Willamette	256	100	15.0X	478
C 2.0	2000	Northwood	128	100	20.0X	478
P4 1.6	1600	Willamette	256	100	16.0X	423
P4 1.6	1600	Willamette	256	100	16.0X	478
C 2.1	2100	Northwood	128	100	21.0X	478
P4 1.7	1700	Willamette	256	100	17.0X	423
P4 1.7	1700	Willamette	256	100	17.0X	478
C 2.2	2200	Northwood	128	100	22.0X	478
P4 1.6	1600	Northwood	512	100	16.0X	478
C 2.3	2300	Northwood	128	100	23.0X	478
C 2.4	2400	Northwood	128	100	24.0X	478
C 2.5	2500	Northwood	128	100	25.0X	478
P4 1.8	1800	Northwood	512	100	18.0X	478
C 2.6	2600	Northwood	128	100	26.0X	478
C 2.7	2700	Northwood	128	100	27.0X	478
C 2.8	2800	Northwood	128	100	28.0X	478
P4 2.0	2000	Northwood	512	100	20.0X	478
P4 2.2	2200	Northwood	512	100	22.0X	478
C D 320	2400	Prescott	256	133.3	18.0X	478
P4 2.4	2400	Northwood	512	100	24.0X	478
C D 325	2533	Prescott	256	133.3	19.0X	478
C D 325/J	2533	Prescott	256	133.3	19.0X	T/775
P4 2.26B	2267	Northwood	512	133.3	17.0X	478
C D 330	2667	Prescott	256	133.3	20.0X	478
C D 330/J	2667	Prescott	256	133.3	20.0X	T/775
P4 2.4B	2400	Northwood	512	133.3	18.0X	478
P4 2.6	2600	Northwood	512	100	26.0X	478
P4 2.4A*	2400	Prescott	1024	133.3	18.0X	478
C D 335	2800	Prescott	256	133.3	21.0X	478
C D 335/J	2800	Prescott	256	133.3	21.0X	T/775
P4 2.53B	2533	Northwood	512	133.3	19.0X	478
C D 340	2933	Prescott	256	133.3	22.0X	478
C D 340/J	2933	Prescott	256	133.3	22.0X	T/775
P4 2.4C	2400	Northwood	512	200	12.0X	478
P4 2.66B	2667	Northwood	512	133.3	20.0X	478
P4 2.8B	2800	Northwood	512	133.3	21.0X	478
P4 2.6C	2600	Northwood	512	200	13.0X	478
P4 2.8A*	2800	Prescott	1024	133.3	21.0X	478
P4 2.8E	2800	Prescott	1024	200	14.0X	478
P4 520/J	2800	Prescott	1024	200	14.0X	T/775
P4 3.06B HTT	3067	Northwood	512	133.3	23.0X	478
P4 2.8C	2800	Northwood	512	200	14.0X	478
P4 3.0E	3000	Prescott	1024	200	15.0X	478
P4 530/J	3000	Prescott	1024	200	15.0X	T/775
P4 3.0C	3000	Northwood	512	200	15.0X	478
P4 3.2E	3200	Prescott	1024	200	16.0X	478
P4 3.2C	3200	Northwood	512	200	16.0X	478
P4 3.4E	3400	Prescott	1024	200	17.0X	478
P4 550/J	3400	Prescott	1024	200	17.0X	T/775
P4 3.4C	3400	Northwood	512	200	17.0X	478
P4 580/J	3600	Prescott	1024	200	18.0X	T/775
P4EE 3.2	3200	Gallatin	512	200	16.0X	478
P4 540/J	3800	Prescott	1024	200	19.0X	T/775
P4 570J	3800	Prescott	1024	200	19.0X	T/775
P4EE 3.4	3400	Gallatin	512	200	17.0X	478
P4EE 3.4	3400	Gallatin	512	200	17.0X	T/775
P4 580J	4000	Prescott	1024	200	20.0X	T/775
P4EE 3.46	3467	Gallatin	512	266	13.0X	T/775
P4EE 3.73	3733	Prescott	2048	266	14.0X	T/775

* Prescott 2.4A and 2.8A processors have HyperThreading Technology (HTT) disabled.

*** Front Side Bus (FSB) Speeds are "quad pumped", so Intel's FSB numbers are four times the actual bus speed on Pentium 4/Celeron, Pentium M/Celeron M, and Itanium processors. Multipliers are based off the base bus speed, not the FSB value.

NetBurst consists of a deep 20-stage pipeline coupled to an 8-stage fetch/decode unit. Due to the time spent fetching and decoding instructions, Intel created a new type of cache called a trace cache. This contained pre-decoded micro-ops, so for a large percentage of instructions, NetBurst runs as a 20-stage pipeline. Certain types of code run very well on NetBurst, while others - specifically branch-heavy code, like that seen in compilers and some games - do not. An incorrect branch prediction on P4 costs about twice as many lost cycles as an incorrect branch prediction on P3 or Athlon, which is why Intel added a more robust branch prediction unit.

The long pipeline allowed clockspeeds to scale very quickly with NetBurst. It was also a bandwidth hungry design, so increasing bus speeds combined with dual-channel memory eventually pushed the P4 beyond the reach of the Athlon XP. On the server front with the Xeon processors, the bandwidth was provided by adding L3 cache.

The Prescott further entered the NetBurst pipeline to 31 stages in addition to the 8 fetch/decode stages. For relative comparison, Intel generally describes the pipeline of the Prescott as 34 stages while also calling the

The Prescott further extended the NetBurst pipeline to 23 stages in addition to the 6 fetch/decode stages. For whatever reason, Intel generally describes the pipeline of the Prescott as 31 stages while only calling the earlier design a 20 stage pipeline. Besides the additional stages, Prescott doubled the L2 cache of the Northwood, added SSE3 support, and to the best of my knowledge contains deactivated x86-64 support - called EM64T by Intel and AMD64 by its creator AMD. Xeon versions of Prescott with the 64-bit support enabled are now shipping, and likely by the time XP-64 is released we will see 64-bit enabled desktop processors.

The Pentium 4 architecture also saw the introduction of Symmetric Multi-Threading (SMT) for Intel processors - they chose to call it Hyper Threading Technology (HTT). It appears to have been a part of the core from the very beginning, but Intel didn't enable the functionality until the P4 3.06 was launched, at which time it became available in the Xeon platforms as well. Later, it was enabled in all the 800 FSB "C" processors. Due to the length of the P4 pipeline, HTT allows the execution units to stay busy in the event of an incorrect branch prediction. The second thread can continue to run while the other thread recovers. In an ideal scenario, HTT could potentially increase performance by 20 or even 50 percent. In real world tests, however, rarely does it improve performance by more than 5 to 10 percent, and there are even times when it hurts performance.

With the switch to socket 775 LGA, Intel has also adopted model names. This likely has something to do with the recent difficulties Intel has encountered in scaling the NetBurst architecture to higher speeds. However, an even bigger problem is Intel's own Pentium M architecture (which is the next section). Anyway, we now have new model numbers which are supposed to reflect the overall capabilities of the chip, with higher numbers indicating more desirable chips. Comparing between families of chips should not be done based solely off the model number, however - there will certainly be instances where a 5xx chip offers better performance than a 7xx chip, and perhaps we'll also see some 3xx chips outperform their "superiors". For the time being, all of the 5xx chips are Prescott cores with 1 MB of L2 cache and an 800 MHz FSB. Future processors are also listed, and you can see where they will likely fall in the performance spectrum.

Mobile Celeron, Mobile P4, Celeron M and Pentium M Processors

						Multiplier
Mobile Pentium/Celeron Chips**						
MC 1.4	1400	Willamette	128	100	14.0X	478M
MC 1.5	1500	Willamette	128	100	15.0X	478M
MC 1.6	1600	Willamette	128	100	16.0X	478M
MC 1.7	1700	Willamette	128	100	17.0X	478M
MC 1.4	1400	Northwood	256	100	14.0X	478M
MC 1.8	1800	Willamette	128	100	18.0X	478M
MC 1.5	1500	Northwood	256	100	15.0X	478M
MC 2.0	2000	Willamette	128	100	20.0X	478M
MC 1.6	1600	Northwood	256	100	16.0X	478M
CM 353/J	900	Dothan	1024	100	9.0X	478M
MC 2.1	2100	Willamette	128	100	21.0X	478M
CM 333	900	Banias	1024	100	9.0X	478M
PM 900 (ULV)	900	Banias	1024	100	9.0X	478M
MC 1.7	1700	Northwood	256	100	17.0X	478M
MC 2.2	2200	Willamette	128	100	22.0X	478M
MC 1.8	1800	Northwood	256	100	18.0X	478M
CM 373/J	1000	Dothan	1024	100	10.0X	478M
MC 2.3	2300	Willamette	128	100	23.0X	478M
PM 1.0 (ULV)	1000	Banias	1024	100	10.0X	478M
MC 2.4	2400	Willamette	128	100	24.0X	478M
MC 2.0	2000	Northwood	256	100	20.0X	478M
PM 723/J (ULV)	1000	Dothan	2048	100	10.0X	478M
PM 1.1 (LV)	1100	Banias	1024	100	11.0X	478M
CM 350/J	1300	Dothan	512	100	13.0X	478M
MC 2.2	2200	Northwood	256	100	22.0X	478M
PM 1.2 (LV)	1200	Banias	1024	100	12.0X	478M
CM 320	1300	Banias	512	100	13.0X	478M
MC 2.4	2400	Northwood	256	100	24.0X	478M
PM 1.3	1300	Banias	1024	100	13.0X	478M
PM 718 (LV)	1300	Banias	1024	100	13.0X	478M
CM 330	1400	Banias	512	100	14.0X	478M
MC 2.5	2500	Northwood	256	100	25.0X	478M
CM 360/J	1400	Dothan	1024	100	14.0X	478M
MC 2.6	2600	Northwood	256	100	26.0X	478M
CM 340	1500	Banias	512	100	15.0X	478M
PM 1.4	1400	Banias	1024	100	14.0X	478M
PM 713 (ULV)	1400	Banias	1024	100	14.0X	478M
MC 2.7	2700	Northwood	256	100	27.0X	478M
CM 370/J	1500	Dothan	1024	100	15.0X	478M
MC D 325	2533	Prescott	256	133.3	19.0X	T/775
MC 2.8	2800	Northwood	256	100	28.0X	478M
PM 1.5	1500	Banias	1024	100	15.0X	478M
PM 705	1500	Banias	1024	100	15.0X	478M
PM 733/J (ULV)	1400	Dothan	2048	100	14.0X	478M
PM 738/J (LV)	1400	Dothan	2048	100	14.0X	478M
MC D 330	2667	Prescott	256	133.3	20.0X	T/775
MC D 335	2800	Prescott	256	133.3	21.0X	T/775
PM 1.6	1600	Banias	1024	100	16.0X	478M
PM 715	1500	Dothan	2048	100	15.0X	478M
PM 758/J (LV)	1500	Dothan	2048	100	15.0X	478M
MC D 340	2933	Prescott	256	133.3	22.0X	T/775
PM 1.7	1700	Banias	1024	100	17.0X	478M
MC D 345	3066	Prescott	256	133.3	23.0X	T/775
MP4 2.8	2800	Northwood	512	133.3	21.0X	478M
MP4 2.8 HT	2800	Northwood	512	133.3	21.0X	478M
PM 735	1700	Dothan	2048	100	17.0X	478M
MC D 350	3200	Prescott	256	133.3	24.0X	T/775
PM 730/J	1600	Dothan	2048	133.3	12.0X	478M
MP4 518	2800	Prescott	1024	133.3	21.0X	?478M
PM 745	1800	Dothan	2048	100	18.0X	478M
PM 753/J (ULV)	1800	Dothan	2048	100	18.0X	478M
MP4 3.0	3000	Northwood	512	133.3	22.5X	478M
MP4 3.0 HT	3000	Northwood	512	133.3	22.5X	478M
PM 740/J	1733	Dothan	2048	133.3	13.0X	478M
MP4 532	3067	Prescott	1024	133.3	23.0X	?478M
MP4 3.2 HT	3200	Northwood	512	133.3	24.0X	478M
MP4 538	3200	Prescott	1024	133.3	24.0X	?478M
PM 750/J	1867	Dothan	2048	133.3	14.0X	478M
PM 755	2000	Dothan	2048	100	20.0X	478M
PM 760/J	2000	Dothan	2048	133.3	15.0X	478M
MP4 552	3467	Prescott	1024	133.3	26.0X	?478M
MP4 558	3600	Prescott	1024	133.3	27.0X	?478M
PM 770/J	2133	Dothan	2048	133.3	16.0X	478M
PM 765	2400	Dothan	2048	100	24.0X	478M

** There are several chips in the mobile sector. PM is for Pentium M, MP4 is the Mobile Pentium 4, CM is the Celeron M, and MC is the Mobile Celeron (P4 core).

*** Front Side Bus (FSB) Speeds are "quad pumped", so Intel's FSB numbers are four times the actual bus speed on Pentium 4/Celeron, Pentium M/Celeron M, and Itanium processors. Multipliers are based off the base bus speed, not the FSB value.

With the scaling clock speeds of the Pentium 4, not even the specially designed Mobile versions were really suited for use in laptops. (Of course, they were still used, but Intel had other plans.) Higher clockspeeds mean higher power requirements as well as increased heat output, which makes it very difficult to get increased battery life. In response to pressure from companies such as Transmeta, Intel commissioned a design team in Israel to put together a high-performance, low-power processor. The end result was the Pentium M. Where the push for high clockspeeds was the driving force behind the NetBurst design, Pentium M was targeted at reaching specific thermal requirements. While specific details are rather hard to come by, since Intel is trying to protect its lead in the Mobile space, the Pentium M appears to be a modified version of the venerable P6 architecture.

One of the improvements made to the P6 architecture was a large L2 cache, which could be powered and cooled in 2W sections. This allows large sections of the cache to be in a low power "sleep" mode at any

One of the improvements made to the P4 architecture was a large L2 cache, which could be powered and accessed in 32K sections. This allows large portions of the cache to be in a low-power "sleep" mode at any given time, so they get the performance benefit of a large cache without incurring as much of the usual power increase. The L1 cache was also doubled from the PIII to 32K+32K data and instruction. Floating point performance was increased with the doubling of MMX/SSE units - although this really only helped with SSE optimized code - and there were a few other architectural changes. Overall, the Pentium M is able to provide performance that's roughly the equivalent of an Athlon processor of the same clock speed, while requiring much less power. Battery life in laptops that use the Pentium M can often be 25 to 50 percent longer than equivalent laptops that use the Mobile Pentium 4, Mobile Celeron or Mobile Athlon XP chips.

The length of the above chart should be an indication of how big the mobile market has become. One of the reasons for this increase in size is likely the cut-throat conditions that exist in the desktop CPU market. Intel charges a hefty premium for most of their mobile processors since, generally speaking, anyone looking for a high-performance laptop has more money to burn. This is what I call the "mobility tax": you should only buy a laptop if portability is a *primary* concern; otherwise, your money will go a lot further with a desktop system. Certainly, business types that use computers for presentations and work on the road will be willing to pay this so-called tax.

With the release of the Dothan core Pentium M chips, Intel has also switched to model numbers. Here, however, there are many factors that influence the overall number. Ultra-Low Voltage processors running at lower clock speeds can end up rated higher than faster processors that require more power. This is supposed to reflect the relative desirability of certain features, as an increased battery life could be more important to some people than raw performance. Of course, Intel specifically states that the model numbers are not measures of performance, but only the technically literate are likely to know this. In their own words: "Intel processor numbers are not a measure of performance. Processor numbers differentiate features within each processor family, not across different processor families. See http://www.intel.com/products/processor_number for details."

Itanium and Itanium 2 Processors

					Multiplier	L3 Cache	
Itanium (Server)							
Itanium	733	Merced	96	66	11.0X	PAC-418	2048
Itanium	733	Merced	96	66	11.0X	PAC-418	4096
Itanium	200	Merced	96	66	12.0X	PAC-418	2048
Itanium	200	Merced	96	66	12.0X	PAC-418	4096
Itanium 2	900	McKinley	256	100	9.0X	PAC-611	1536
Itanium 2	900	McKinley	256	100	9.0X	PAC-611	3072
Itanium 2	1000	McKinley	256	100	10.0X	PAC-611	1536
Itanium 2	1000	McKinley	256	100	10.0X	PAC-611	3072
Itanium 2 LV	1000	Deerfield	256	100	10.0X	PAC-611	1536
Itanium 2 LV	1500	Deerfield	256	100	15.0X	PAC-611	1536
Itanium 2	1300	Madison	256	100	13.0X	PAC-611	3072
Itanium 2	1400	Madison	256	100	14.0X	PAC-611	4096
Itanium 2	1500	Madison	256	100	15.0X	PAC-611	6144

*** Front Side Bus (FSB) Speeds are "quad pumped", so Intel's FSB numbers are four times the actual bus speed on Pentium 4/Celeron, Pentium M/Celeron M, and Itanium processors. Multipliers are based off the base bus speed, not the FSB value.

Itanium processors are likely one of the least understood CPUs by most computer enthusiasts. Given that the cheapest models still cost over \$1000, that's not really surprising. These processors are meant to target the high-end corporate world. They are often used in massively parallel processing situations, and Itaniums are capable of working in up to 512-way SMP systems. Of course, that doesn't really explain *what* the Itanium is.

For starters, Itanium is the way that Intel envisioned 64-bit computing, and it is built on a new instruction set dubbed IA-64 (Intel Architecture 64).. IA-64 was a clean break from x86 legacy code, and it was designed for the future. Really, its competition isn't the Xeon or Opteron CPUs, although some mistakenly compare it with these processors. Itanium is meant to compete in the high-end corporate 64-bit computing world, going up against servers based on the IBM Power4/5, HP PA-RISC, Sun UltraSparc-III, and DEC Alpha. If none of those names ring a bell, that's not very surprising. The quad-processor IBM Power4 system that was used as the main server at a company I worked for (and they had two units for redundancy) cost somewhere in the neighborhood of \$500,000, and the RAID-5 array that provided data storage was another \$500,000. Perhaps more important than the hardware was the service contract with IBM that helped guarantee everything stayed running. The cost of the support contract with IBM (for dozens of such setups) was supposedly around \$300 million dollars a year!

The Alpha technology, interestingly enough, was purchased by Compaq, who merged with HP, and HP worked with Intel on the design of the Itanium, with the intention of using it in place of PA-RISC once it was complete. I believe that some (all?) of the Alpha technology was later transferred to Intel, most likely for use in furthering the design of the Itanium processors. Compaq/HP has continued to support this chip for the past several years, but they haven't invested a lot of money into researching new iterations of the design. This makes sense, since HP is encouraging its enterprise customers to switch to their Itanium platforms. Recently, HP announced that the 1.3 GHz (I think that was the speed) EV7 version of the Alpha chip will be the last.

These systems are often referred to as "Big Tin" systems, and they're in a league of their own. They are frequently used in systems that process *huge* amounts of data - their 64-bit addressing allows the use of many gigabytes of physical RAM - and they are usually optimized for input/output functions. Of course, reliability and up-time are far more important than actual performance numbers, and often once a system has been built around a specific architecture, large corporations will stick with that hardware unless there is *tremendous* incentive to switch to something else. Switching usually consists of several years of coding, testing, debugging, and validation - a task not to be undertaken lightly, to be sure.

For the processor design, Intel continued with their radical departure from accepted norms. Instead of a RISC or CISC approach, Intel went back to a technology that had been used in old mainframes and other computers of yore. VLIW (Very Long Instruction Word). Itanium is not a strict VLIW machine, though, as VLIW has some well known drawbacks that Intel worked to overcome, and Intel chose to call their new approach EPIC, "Explicitly Parallel Instruction Computer". In contrast to designs such as the Xeon and Opteron, which can issue up to three instructions per cycle, the Itanium 2 (forget Itanium 1 for a minute) can issue *eight* instructions per clock, and unlike VLIW designs, future Itanium chips could further increase the issue width without needing to recompile the code. In theory, then, a 1 GHz Itanium chip could perform roughly as fast as a 2.66 GHz Xeon/Opteron, or the 1.5 GHz Itanium 2 would be roughly as fast as a 4 GHz Xeon/Opteron. That's just theoretical performance, of course, and the overall system design will play a large role in determining how much of the potential of any system is actually realized.

To help reach that potential, Itanium chips run off a 128-bit quad-pumped system bus, using standard SDRAM (for the time being). The lower clock speeds combined with the wider bus make the SDRAM less of an issue than with high-speed desktop systems. The initial Itanium design, Merced, had four integer units (ALUs), two floating point units (FPU's), and three branch units (BRUs), two SIMD (i.e. MMX/SSE) units, and two load/store units - also called address generation units (AGUs) in other CPUs. The modified McKinley (and later) designs have six ALUs, three BRUs, two FPU's, one SIMD, two load units, and two store units - sort of like having 4 AGUs, except that they're more specialized. In addition, the McKinley has roughly three times the cache bandwidth as Merced. Merced was also a six issue design with a deeper pipeline (10 stages) and less memory bandwidth - a rather problematic design. McKinley and later designs are eight issue designs with shorter pipelines (8 stages) and more memory bandwidth. While Merced rarely made full use of its six issue design, McKinley's enhancements help it come a lot closer to issuing the maximum eight instructions per clock.

That doesn't really tell a whole lot about the architecture, and I don't really want to go much deeper than that right now. Suffice it to say that Itanium depends in a large part on compiler technology in order to reach its potential, and Intel has apparently had more difficulties in that area than they initially anticipated, but lately this seems to be less of a problem. The initial Merced design was also flawed, if you couldn't tell from my above description of the architecture, but Itanium 2 goes a long way toward rectifying the problems.

Many have called the Itanium a failure - coming up with such names as Itanic to describe the processor - especially now that AMD has launched Opteron and Intel is following suit with x86-64 support. However, they're really very different goals, and in the target market segment, Itanium is still managing to compete. Needless to say, it helps that Intel has very deep pockets thanks to the income generated from their desktop and mobile processor divisions. Itanium may or may not live in the long term, but short term Intel has plans to keep it around at least another three or four years, and they will likely keep it around longer to support existing clients. Honestly, though, I doubt any of us will ever be running an IA64 processor on our desktop systems.

Conclusion

Obviously, that's only a brief glimpse at the processor histories of AMD and Intel, with a vague picture of the future. Dual core designs should start appearing within the next year, and rumors of quad core processors are also floating around the web. At some point, we will likely reach the limits of current manufacturing technologies, but that day is still a long ways off. AMD and Intel both have technologies in development that should carry us past 45 nm process technologies, and probably down to single digits in our lifetime. That's assuming we don't get quantum computers first, that make all of the current binary systems seem quaint by comparison.

The amount of processing power sitting in front of you right now was beyond comprehension a couple decades ago. Even the "average" computers of today would seem amazing to people even one decade in the past. Ten years ago, 3D was only dreamt about, and professional 3D accelerators cost thousands of dollars while doing far less than a "cheap" GeForce 3 or Radeon 8500. Ten years ago, 32-bit processors were still looking for a real operating system, and 64-bit was only used by governments and research centers. Ten years ago, a 100 MHz processor was as good as it got. Ten years ago, few people had ever used a networked computer at home, and 28.8 modems were amazingly fast. Here's hoping the gurus at AMD, Intel, and other companies can continue to amaze us for another ten years!

Stay tuned for more insider articles from Jarred, including a much anticipated GPU cheat sheet as well!

[https://web.archive.org/web/20040909115936/http://www.anandtech.com/cpuchipsets/showdoc.aspx?
i=2178](https://web.archive.org/web/20040909115936/http://www.anandtech.com/cpuchipsets/showdoc.aspx?i=2178)

ANANDTECH

your source for hardware analysis and news

- Motherboards
- Memory
- Storage
- Cases/Cooling
- IT Computing
- Displays
- Mobile
- Macintosh
- CPU/Chipsets
- Video
- Digital Cameras
- Linux
- Multimedia
- Systems
- Trade Shows
- Guides

- Home
- About
- Forums
- RSS/Mobile Users
- News
- RealTime Pricing
- My AnandTech
- Team AnandTech
- Anand's Weblog

Search AnandTech

GO!

Huge August Rebates
at
Shop.Kingston.com!
Click here for details!

Dear ZipZoomFly Shoppers,

To better server you, our website will be unavailable for an
estimated time of 4 hours for system maintenance. We apologize.

Updated CPU Cheatsheet - Seven Years of Covert CPU Operations



Date: August 28th, 2004

Topic: CPU & Chipset

Manufacturer: Various

Author: Jarred Walton

AnandTech Deals

AMD Athlon 64 3000+ 2.048 GHz (ada3000box) Processor Products

Prices

- ZipZoomFly.com \$170.00
- Amazon \$254.99
- PC Connection \$224.84

Introduction

Update: 8/27/04 - The charts have all been revised. Thanks go out to all the people that posted corrections in the comments section as well as sending them via email. In addition to the corrections, some further information and commentary has been added to the pages. For anyone that actually comes back to this article for reference information, enjoy the changes!

Foreword by Kristopher Kubicki:

From time to time we stumble upon some truly gifted and patient people here at AnandTech. Some weeks ago I wrote a [CPU codename cheatsheet](#) as just something to do in an airport terminal to kill time. Very soon after, an extremely diligent Jarred Walton showed me his rendition of the CPU family tree that he was keeping just for fun? Knowing I was bested, I offered Jarred a chance at writing a pilot for AT, and here it is! Please enjoy the second, extremely thorough CPU Cheatsheet 2.0.

*But loud! what lurks in yonder chassis, hot?
A CPU, my programs it will run!
O Pentium, Pentium! wherefore art thou Pentium?
Obscure thy benchmarks and refuse thy name.
What's in a name? that which we call a chip
By any other name would run as fast.*

My sincere apologies to Shakespeare, but that mangled version of *Romeo and Juliet* is an apt description of the world of computer processors. Once upon a time, we dealt with part numbers and megahertz. Larger numbers meant you had a faster computer. 80286 was faster than 8088 and 8086, and the 80386 was faster still, with the 80486 being the king of performance. Life was simple, and life was good. But that is the distant past; welcome to the present.

Where once we had a relatively small number of processor parts to choose from, we are now inundated with product names, model numbers, code names, and features. Keeping track of what each one means is becoming a rather daunting task. Sure, you can always try Googling the information, but sometimes you'll get conflicting information, or unrelated web sites, or only small tidbits of what you're trying to find out. So, why not put together a clear, concise document that contains all of the relevant information? Easier said than done; however, that is exactly what is attempted in this article.

In order to keep things even remotely concise, the cutoff line has been arbitrarily set to the Pentium II and later Intel processors, and the Athlon and later AMD processors. Anything before that might be interesting for those looking at the history of processors, but for all practical purposes, CPUs that old are no longer worth using. Also absent will be figures for power draw and heat dissipation, mainly because I'm not overly concerned with those values, not to mention that AMD and Intel have very different ways of reporting this information. Besides, Intel and AMD design and test their CPUs with a variety of heatsinks, motherboards, and other components to ensure that everything runs properly, so if you use the proper components, you should be fine.

So what will be included? For this first installment, details on clock frequencies, bus speeds, cache sizes, transistor counts, code names, and a few other items has been compiled. The use of model numbers with processors is also something people will likely have trouble keeping straight, so the details of processors for all Athlon XP and later AMD chips and Pentium 4 and later Intel chips will follow. The code names and features will be presented first, with individual processor specifics listed on the later pages. As a whole, it should be a useful quick reference - or cheat sheet, if you prefer - for anyone trying to find details on a modern x86 processor.

With that said, on to the AMD processors. Why AMD first? Because someone has to be first, and AMD comes before Intel in the alphabet.

AMD Processors >

[Click Here](#)

Review Index:

Select

Print this article

Find the lowest prices

Email this article

Buy it from ZipZoomFly.com for \$170.00

70 comment(s) - last comment [Jenand](#) on Sep 8, 2004 at 11:24 AM

Latest News

August 26, 2004

Dual-core chips bring dual caches

July 28, 2004

AMD Sempron Reviews

July 27, 2004

More AMD CPU price changes hit the market

Itanium, Xeon to share common system platform

IBM to produce a dual-core G5

July 10, 2004

AMD A64 90nm Benchmarks and Info

July 9, 2004

64-Bit Battle : Intel's Xeon 3.4 GHz vs. AMD's Opteron 250

May 18, 2004

AMD's Opteron 150 and 250 processors

April 30, 2004

Kicking It Down A Notch : AMD's Athlon64 2800+ Processor

April 27, 2004

64-bit vs 32-bit Performance [Athlon 64]

April 12, 2004

AMD Athlon 64 2800+ Processor Review

March 30, 2004

Intel Celeron 2.8GHz Review

March 22, 2004

AMD Athlon 64 3000+ Processor Performance

Intel Pentium 4 3.4E Reviews

March 8, 2004

AMD Athlon 64 Model 3000+ Review

February 23, 2004

AMD 64 3400+ vs FX-51 @ Bjorn3D

February 18, 2004

The AMD K8 Architecture

more News

Anand's Picks

AMD Sempron 3100, 1.8 GHz (sda3100axbox) Processor - \$130

EPoX EP-8KDA3+ Motherboard - \$142

Hitachi Deskstar 7K400 (0a30228) Hard Drive - \$449.37

Abit AV8 Motherboard - \$125

AMD Athlon 64 3000+, 2.048 GHz (ada3000box) Processor - \$170

Anand's Blog

Getting Married

What a week

Ok so it's not a Macdate

Latest CPU & Chipset Forum Posts

Barton magic steppings

mobile celeron VS mobile pentium III

4000+ Only 2.4Ghz !!!!!

Sponsored Links (Get Listed)

[Best deals on CPU's only @ TigerDirect.com](#)

We have what you are looking for in stock & ready to ship at a great price! Computers, Notebooks, Monitors, Digital Cameras, CPU's, Memory, Motherboards, Barebone kits, Networking, Video, & more at the best values. Order today, ships today!

[Hard-to-Find Computer Parts from StarTech.com](#)

StarTech.com is the professionals' source for hard-to-find computer parts. Try our online Parts Finder tool to find the product you need quickly.

[Priority Electronics- Notebook Online Store](#)

Online seller of notebook batteries, accessories, and parts, featuring low prices for all major brands, such as, Dell, Sony, IBM, Toshiba, Compaq, HP, Gateway and more.

[Velocity Micro Gaming Computers](#)

Performance matters. Velocity Micro proves a computer is more than the sum of it's components. Experience our powerful Velocity Generator configuration tools to let us build your dream system!

[Aberdeen - Custom Rackmount Storage, Servers & NAS](#)

Aberdeen, an innovator of storage solutions provides custom computing solutions such as; XDAS Network Attached Storage servers, 1U to 5U rackmount servers, Supermicro Dual Xeon servers, chassis & motherboards and HPC Grid computing solutions.

Another J400+ question...this time a little different.

Desirable CPUs drying up?

DFI Lanparty UT NF3 250GB review on AT!

P4 3.0GHz Overclock to 3.22GHz with stock fan/heatsink no problem?

HELP Overclocking a AMD 64 3200+ with a K8N-e Deluxe

OC my A64 2800... Have a Q.

P4M On A P4 Laptop ?

[more CPU & Chipset Discussions](#)

Copyright © 1997-2004 AnandTech, Inc. All rights reserved. Terms, Conditions and privacy information.
[Click Here for Advertising Information](#)




EXHIBIT B

<https://web.archive.org/web/20121206115711/http://www.ece.ucdavis.edu/vcl/pubs/2011.12.techreport.techscale/techscaling.pdf>

Toward More Accurate Scaling Estimates of CMOS Circuits from 180 nm to 22 nm

Aaron Stillmaker, Zhibin Xiao, and Bevan Baas

VLSI Computation Lab

Department of Electrical and Computer Engineering

University of California, Davis

Technical Report ECE-VCL-2011-4

{astillmaker, zxiao, bbaas}@ucdavis.edu

Abstract—With deep submicron technology nodes other methods are needed to obtain scaling factors rather than the traditional scaling factors which held for the pre-submicron era. This work presents scaling factors between major technology nodes between 180 nm and 22 nm operating at voltages from 1.8 V to 0.7 V. Common operating data for these technologies were taken from the International Technology Roadmap for Semiconductors (ITRS). HSpice simulations that rely on the Predictive Technology Model (PTM) for transistor characteristics were used to find the scaling factors.

I. INTRODUCTION

It is useful to be able to compare between digital design results of circuits that are made using different technology sizes. It has become apparent that traditional scaling does not hold with deep submicron technology sizes. This work models devices from different technology sizes to get the scaling factors of energy, delay, and area between sizes. One of the large motivations for this project was the ability to compare different chips in a fair fashion, the size and voltage were varied over a large range of values, and the power, energy, and intrinsic delay were all measured. From these numbers, scaling factors were extrapolated.

A. Traditional Scaling Methods

Until the deep submicron era, transistor characteristics scaled very predictably as designers generally focused on simple geometry scaling. This included generalized equations which take short-channel effects into consideration. With this, great gains were to be had following this method, with few outside factors forcing designers to optimize non-traditionally. These scaling factors have been used and taught, so they are easily found in the literature [1], [2], and is shown in Table I where scaling factor S is the ratio of the all transistor geometry between two transistor sizes and U is the ratio between two voltages. With both S and U it is expected that all geometry and voltages scale together.

B. Scaling Equations in Submicron Technologies

As transistors get smaller, short channel effects and other issues such as process variation start playing a larger role, making the traditional scaling equations inaccurate. Leakage current is affected greatly by gate length, oxide thickness, and

threshold voltage, so it is becoming a large issue with deep submicron processes. With these new issues now affecting transistor operation, designers started looking to optimize between technology nodes other than simple geometric scaling. Width, length, and oxide thickness are not scaling together as well as V_{DD} and V_T . Which means that scaling factors S and U shown in Table I can not be determined. The above mentioned problems are especially noticeable when the industry switched largely to using high- κ dielectrics and metal gates with technology nodes at 45 nm and smaller.

C. Methods for Accurate Scaling

The physics that affect transistors as they reach the submicron region get far more complicated, with leakage, and other issues becoming a large factor in energy consumption and delay. So the most accurate way to get scaling factors in submicron processes is to use a simulation tool, such as HSpice with a model that specifies the characteristics of the particular technology. Simulating a whole design in Spice with modified technology size and voltages would result in the most accurate comparison, but this is not a practical solution because other issues such as transistor sizing can arise with a large change. Not to mention, one generally does not have the exact specifications of the design they want to compare with. Therefore a good approximation of a device in a certain technology is to use inverters in a chain, with 4 inverters attached to each output, this is known as Fan Out 4, or FO4. A circuit that has a delay and consumption of X number of FO4 inverter chains in a certain technology size, should have roughly the same X number of FO4 inverter chains in a different technology size. With this in mind, this work sets out to take measurements of FO4 models in many different sizes and voltages so that roughly accurate scaling factors can be used.

II. INTERNATIONAL TECHNOLOGY ROADMAP FOR SEMICONDUCTORS (ITRS)

The International Technology Roadmap for Semiconductors (ITRS) [3], [4], [5], [6], [7], [8] is a report which predicts where semiconductor technology is headed in the next 15 years, and is made on odd years, with updates on even

TABLE I
SCALING SCENARIOS FOR SHORT-CHANNEL DEVICES,
PRE-SUBMICRON [1].

Parameter	Relation	Full Scaling	General Scaling	Fixed Voltage Scaling
W, L, t_{ox}		$1/S$	$1/S$	$1/S$
V_{DD}, V_T		$1/S$	$1/U$	1
Area/Device	WL	$1/S^2$	$1/S$	$1/S^2$
Power	$I_{sat}V$	$1/S^2$	$1/U^2$	1
Intrinsic Delay	$R_{on}C_{gate}$	$1/S$	$1/S$	$1/S$
Energy	Pt	$1/S^3$	$1/SU^2$	$1/S$

years. These reports are formed by a collaboration of many companies and research institutions. In this work, these were used to get industry standard technology sizes, and voltages commonly used, as well as general knowledge about transistor changes over the years. Area is also of interest in digital design. Standard 1/2 metal 1 sizes were taken to make scaling factors. In this report, when technology process sizes are referred to, i.e. 180 nm, 45 nm, etc. we are referring to the minimum feature size. Process sizes were generally identified by their smallest feature size, and for a long time DRAM 1/2 pitch sizes were the smallest, and therefore used to identify technologies. With new fabrications, this has not been the case, and IRTS no longer identifies technologies by their minimum feature size to try and stop confusion, they differentiate using the first year of production [4]. However, as minimum feature technology notes have continued to be the generally accepted term, and because we are comparing with older technologies, in this work we will be identifying technologies by both the production year, and minimum feature size, as shown in Table II.

III. PREDICTIVE TECHNOLOGY MODEL (PTM)

A. PTM Overview

The Predictive Technology Models [9], [10], [11], or PTM, was used to simulate different design characteristics as technology size, and voltage changed. The models were developed for designers who do not have access to proprietary transistor characteristics to test designs with future technology nodes.

B. HSpice Modeling

HSpice was used to model the scaling results. A fan out four, FO4, inverter chain was used. FO4 delay has been shown to be proportional to CV/I (intrinsic capacitance, voltage, and drive current of a device) [12]. Intrinsic capacitance and drive current are both proportional to device size, so they scale with $1/S$ and as previously mentioned, voltage scales at factor U , thus delay should scale with U/S^2 . The chain starts with one minimum size inverter with the output connected to 4 inverters, with that output connected to 16 inverters, and so on until the circuit ends with 64 inverters, creating a total of 4 FO4 stages. A square wave input was sent into the chain. The set of 16 inverters, directly in the middle of the chain, was used for the sampling. The delay between when the input signal to the set

of 16 inverters crossed the mid point and the output crossed the midpoint was measured. The voltage was measured along with the current, and calculations were made by the simple equations 1-4 where t_0 to t_1 is the transition time as the signal goes high, and t_2 to t_3 is the transition time as the signal goes low.

$$P_{ave} = \frac{1}{T} \int_0^T I(t) \cdot V dt \quad (1)$$

$$E_{up} = \int_{t_0}^{t_1} I(t) \cdot V dt \quad (2)$$

$$E_{down} = \int_{t_2}^{t_3} I(t) \cdot V dt \quad (3)$$

$$E_{ave} = \frac{E_{up} + E_{down}}{2} \quad (4)$$

All these calculations are done by HSpice and provided as a simulation output.

1) *Simulation Parameters:* The simulations were run on the technology sizes: 180 nm, 90 nm, 65 nm, 45 nm, 32 nm, and 22 nm with supply voltages varying from 1.3 V to 0.7 V in 0.1 V increments and 1.8 V. The 180 nm stopped working at 0.8 V and 0.7 V, as this is a lower voltage than standard 180 nm was expected to run at. Similarly technology nodes are not designed to handle voltages much higher than their target voltages, so even though HSpice gave results for the smaller technology nodes operating at higher voltages, they were removed from the results as the PTM characteristics should not hold for these values. These spots are marked as N/A in the results tables.

With the industry standard of using high- κ dielectric transistors at 45nm and below, high- κ PTM models were used as the most accurate models available because semiconductor companies do not readily provide characteristics of their specific technologies. This level of generality can actually help when being used to compare two chips of different fabrications.

IV. SIMULATION RESULTS AND SCALING FACTORS

As previously mentioned, industry standard numbers were taken from IRTS reports. Table II shows the standard values labeled by IRTS as a high performance circuit with a heat sink, at each technology node that was investigated, as well as the delay and power simulated using the inverter chain in HSpice as described in Section III-B. The V_{DD} is taken from the IRTS tables for high-performance.

A. Area

To determine a factor for scaling area between technologies, minimum feature sizes and Metal 1 half pitches were taken from the IRTS reports. The Logic Gate (4 Transistors) size was taken from the IRTS tables of MPU (High-volume Microprocessor). An exact scaling would be dependent on the design, but using either of the aforementioned values should give a good rough estimate. The relative scaling of area is shown in Figure 1 and detailed in Table III. Tables IV-VII can be used to scale areas depending on preference. To scale find the

TABLE II
CHARACTERISTICS OF DIFFERENT TECHNOLOGIES NODES FOR HIGH
PERFORMANCE PROCESSOR WITH A HEATSINK.

Production Year	Tech. Node (nm)	Min. Vdd (V)	Max. Power (W)	Energy (fJ)	Delay (ps)
1999	180	1.8	90	19	80.0
2001	130	1.2	130	3.4	35.5
2004	90	1.1	149	1.7	27.2
2007	65	1.1	189	1.1	20.4
2008	45	1.1	146	0.64	7.82
2010	32	0.97	146	0.30	6.33
2012	22	0.9	158	0.17	5.12

TABLE III
AREAS OF DIFFERENT ASPECTS OF DIFFERENT TECHNOLOGY NODES [3],
[4], [5], [6], [7], [8].

Minimum Feature Size (nm)	Metal I Half Pitch (nm)	(4T) Logic Gate Size (μm^2)	Relative Area from Feature Size	Relative Area from M1 $\frac{1}{2}$ Pitch	Relative Area from Logic Gate
180	230	57	7.7	11	22
130	150	10.4	4	4.9	4
90	90	5.2	1.9	1.8	2
65	68	2.6	1	1	1
45	59	2.1	0.48	0.75	0.81
32	45	0.71	0.25	0.44	0.27
22	32	0.35	0.11	0.22	0.14

TABLE IV
AREA SCALING USING MINIMUM FEATURE SIZE.

Technology Nodes (nm)	180	130	90	65	45	32	22
180	1	1.9	4.0	7.7	16	32	67
130	0.52	1	2.1	4.0	8.4	17	35
90	0.25	0.48	1	1.9	4.0	7.9	17
65	0.13	0.25	0.52	1	2.1	4.1	8.7
45	0.063	0.12	0.25	0.48	1	2.0	4.2
32	0.032	0.061	0.13	0.24	0.51	1	2.1
22	0.015	0.029	0.060	0.11	0.24	0.47	1

starting technology node from the top column, then go down until the desired technology is reached, then that number can be multiplied by the chip's area to determine an equivalent area in the different technology node.

B. Delay

Table VIII contains the results from the simulations, showing the time for a signal to propagate through one minimum sized inverter in the middle of the FO4 inverter chain. Equation 5 can be used to scale delay where the factor Delay Factor (DF) is found in table VIII.

$$D_x = \frac{DF_x}{DF_y} \cdot D_y \quad (5)$$

(DF from Table VIII)

TABLE V
AREA SCALING USING METAL I HALF PITCH SIZE.

Technology Nodes (nm)	180	130	90	65	45	32	22
180	1	2.4	6.5	11	15	26	52
130	0.43	1	2.8	4.9	6.5	11	22
90	0.15	0.36	1	1.8	2.3	4.0	7.9
65	0.087	0.21	0.57	1	1.3	2.3	4.5
45	0.066	0.15	0.43	0.75	1	1.7	3.4
32	0.038	0.090	0.25	0.44	0.58	1	2.0
22	0.019	0.046	0.13	0.22	0.29	0.51	1

TABLE VI
AREA SCALING USING 4 TRANSISTOR LOGIC SIZE.

Technology Nodes (nm)	180	130	90	65	45	32	22
180	1	5.5	11	22	27	80	160
130	0.18	1	2	4	5	15	30
90	0.091	0.5	1	2	2.5	7.3	15
65	0.045	0.25	0.5	1	1.2	3.7	7.4
45	0.037	0.2	0.4	0.81	1	3	6
32	0.013	0.068	0.15	0.27	0.35	1	2
22	0.0061	0.033	0.068	0.13	0.17	0.49	1

TABLE VII
AREA SCALING FACTORS USING GEOMETRIC MEAN OF THREE ASPECTS
FROM TABLE III.

Technology Nodes (nm)	180	130	90	65	45	32	22
180	1	2.9	6.6	12	19	40	83
130	0.34	1	2.3	4.3	6.4	14	28
90	0.15	0.44	1	1.9	2.8	6.1	13
65	0.080	0.23	0.53	1	1.5	3.3	6.6
45	0.053	0.16	0.35	0.66	1	2.2	4.4
32	0.025	0.072	0.16	0.31	0.46	1	2.1
22	0.012	0.035	0.080	0.15	0.23	0.49	1

Figure 4 shows the results of simulations using the values in Table II compared to what we would expect the results to be if we were using the equations from Table I. Figures 5 shows the scaling of delay with varied technology sizes and operating voltages.

C. Energy and Power

Table IX contains the results from the simulations. Equation 6 can be used to scale energy. Equation 7 combines Equations 5 & 6 using the power equation $Power = Energy/Delay$. The factors Delay Factor (DF) and Energy Factor (EF) are found in table VIII & IX.

$$E_x = \frac{EF_x}{EF_y} \cdot E_y \quad (6)$$

$$P_x = \frac{EF_x \cdot DF_y}{EF_y \cdot DF_x} \cdot P_y \quad (7)$$




Fig. 1. Relative area scaling of different area aspects over different technology nodes.

TABLE VIII

TIME FOR A SIGNAL TO PROPAGATE THROUGH ONE MINIMUM SIZED INVERTER IN THE MIDDLE OF THE FO4 INVERTER CHAIN IN ps.

	1.8 V	1.3 V	1.2 V	1.1 V	1 V	0.9 V	0.8 V	0.7 V
180 nm	80.03	110.3	121.2	134.8	152.2	176.1	N/A	N/A
130 nm	N/A	33.44	35.52	38.29	42.07	47.53	55.84	69.81
90 nm	N/A	N/A	25.2	27.22	30.02	34.1	40.39	51.19
65 nm	N/A	N/A	18.84	20.45	22.7	25.99	31.23	40.49
45 nm	N/A	N/A	7.656	7.824	8.092	8.502	9.166	10.24
32 nm	N/A	N/A	N/A	6.042	6.255	6.576	7.093	7.967
22 nm	N/A	N/A	N/A	N/A	4.878	5.17	5.619	6.359

TABLE IX

ENERGY USED TO TOGGLE A SINGLE MINIMUM SIZED INVERTER IN THE MIDDLE OF THE FO4 INVERTER CHAIN IN fJ.

	1.8 V	1.3 V	1.2 V	1.1 V	1 V	0.9 V	0.8 V	0.7 V
180 nm	19	9.2	7.7	6.3	5.0	3.98	N/A	N/A
130 nm	N/A	4.1	3.4	2.8	2.2	1.8	1.4	1.1
90 nm	N/A	N/A	2.1	1.7	1.4	1.1	0.88	0.69
65 nm	N/A	N/A	1.3	1.1	0.92	0.74	0.59	0.45
45 nm	N/A	N/A	0.80	0.64	0.48	0.36	0.27	0.202
32 nm	N/A	N/A	N/A	0.45	0.33	0.27	0.18	0.13
22 nm	N/A	N/A	N/A	N/A	0.23	0.17	0.12	0.084

(EF from Table IX DF from Table VIII)

Figure 2 shows the results of simulations using the values in Table II compared to what we would expect the results to be if we were using the equations from Table I. Figure 3 show the scaling of energy with varied technology sizes and operating voltages.




Fig. 2. Energy used to toggle one minimum sized inverter in the middle of the FO4 inverter chain simulated using Table II values and scaled using Table I equations.




Fig. 3. Energy required for one minimum sized inverter in the middle of the FO4 inverter chain to toggle for different technology nodes with scaling voltage.

V. CONCLUSION

Comparing the data presented in this report with the traditional scaling methods, indeed does not hold into these submicron transistors. The general trend is similar, but would not make an accurate comparison. Thus the factors gleaned from simulations presented in this work is a more accurate estimation that can be used to compare two devices from different technologies.

VI. ACKNOWLEDGEMENT

The authors gratefully acknowledge support from Fudan University ASIC and System National Lab open project funding, ST Microelectronics, C2S2 Grant 2047.002.014, NSF




Fig. 4. Delay Simulated for a signal to propagate through one minimum sized inverter in the middle of the FO4 inverter chain using Table II values and scaled using Table I equations.




Fig. 5. Delay for a signal to propagate through one minimum sized inverter in the middle of the FO4 inverter chain for different technology nodes with scaling voltage.

Grant 0430090 and CAREER Award 0546907, SRC GRC Grant 1598 and CSR Grant 1659, Intel, UC Micro, Intelliasys, SEM, and a UCD Faculty Research Grant.

REFERENCES

- [1] Jan M. Rabaey, Anantha P. Chandrakasan, and Borivoje Nikolic, *Digital integrated circuits : a design perspective*, Pearson Education, Upper Saddle River, NJ, second edition, 2003.
- [2] John P. Uyemura, *Introduction to VLSI circuits and systems*, John Wiley & Sons, Inc., Hoboken, NJ, first edition, 2002.
- [3] "International technology roadmap for semiconductors 2010 update," Tech. Rep., ITRS, 2010.
- [4] "International technology roadmap for semiconductors 2009 edition," Tech. Rep., ITRS, 2009.

- [5] "International technology roadmap for semiconductors 2007 edition," Tech. Rep., ITRS, 2007.
- [6] "International technology roadmap for semiconductors 2004 update," Tech. Rep., ITRS, 2004.
- [7] "International technology roadmap for semiconductors 2002 update," Tech. Rep., ITRS, 2002.
- [8] "International technology roadmap for semiconductors 2000 update," Tech. Rep., ITRS, 2000.
- [9] W. Zhao and Y. Cao, "New generation of predictive technology model for sub-45nm early design exploration," *IEEE Transactions on Electron Devices*, vol. 53, no. 11, pp. 2816–2823, Nov. 2006.
- [10] Y. Cao, T. Sato, D. Sylvester, M. Orshansky, and C. Hu, "New paradigm of predictive mosfet and interconnect modeling for early circuit design," in *CICC*, 2000, pp. 201–204.
- [11] "Predictive technology model," <http://ptm.asu.edu/>.
- [12] "Fo4 writeup: International technology roadmap for semiconductors 2003 edition," Tech. Rep., ITRS, 2002.

https://web.archive.org/web/20070205153105/http://www.dell.com/downloads/global/vectors/2005_p4_pentiumm.pdf

PERFORMANCE ANALYSIS OF INTEL® PENTIUM® M AND PENTIUM 4 PROCESSORS

Deepak Janakiraman and Chad Roesle, Systems Engineers, Dell™ System Performance Analysis



The increased demand for portable computers with longer battery life, coupled with the growing power budget required for desktop processors, have made it impractical to use desktop processors in mobile platforms. The Intel Pentium M is the first Intel processor designed specifically for mobile platforms. Since its introduction in 2003, it has largely displaced the Mobile Pentium 4 and Pentium 4 processors in portable computers.

The Pentium M architecture balances performance and power consumption. Its architecture allows the Pentium M to operate at lower frequencies, but often outperform its higher-frequency predecessor, the Mobile Pentium 4. In some instances, the Pentium M also performs competitively with the higher-frequency Pentium 4 processors favored in desktop systems.

The Pentium M processor marked a significant shift in the Intel mobile platform architecture. Before the Pentium M, Intel mobile processors were architecturally the same as desktop processors, but were run at lower frequencies to reduce power consumption and meet thermal requirements. For this reason, the performance of these mobile platforms has traditionally been lower than platforms based on higher-frequency desktop processors. With the Pentium M, fundamental architectural changes enable competitive performance at lower frequencies, power consumption, and thermal levels.

This white paper compares the architecture and performance of the Pentium M to higher-frequency Pentium 4 processors. This study highlights the fact that, when evaluating mobile platform performance, factors other than processor frequency are becoming more important.

Processor Architecture

Table 1 compares the processor architecture of the Pentium M and Pentium 4. In order to achieve competitive performance with the thermal constraints of a mobile platform, the Pentium M is equipped with several power-aware performance features: micro-ops fusion, advanced branch prediction, and a dedicated stack manager.

A micro-op is the most basic operation of the microprocessor. Under micro-ops fusion, the processor micro-operations are kept united through most of the out-of-order core of the processor. This effectively reduces the number of micro-ops, which helps to cut down on the energy required to complete the instruction sequence, thereby minimizing power consumption.

Feature	Pentium M	Pentium 4
Transistor process	77 million 90 nanometer (nm)	125 million 90 nm
Die size	84 millimeters (mm)	112 mm
L1 and L2 cache sizes	64 Kilobytes (KB) L1 2 megabytes (MB) L2	16 KB L1 1 MB L2
Bus	533 MHz	800 MHz
Pipeline stages	Less than 31*	31
Features	- Micro-ops fusion - Dedicated stack manager	-
Branch Prediction	Indirect branch predictor and loop detector	Branch predictor
Streaming SIMD Extensions (SSE) support	SSE/SSE2	SSE/SSE2/SSE3
Support for Hyper-Threading Technology (HT)	No	Yes

Source: Intel Corporation

*The shorter pipeline of the Pentium M helps reduce the pipeline-flush recovery time in unpredictable code, which in turn helps improve the overall system performance. In contrast, the longer pipeline of the Pentium 4 allows higher clock speeds, but this comes at the expense of a lower number of instructions per cycle (IPC) and much higher power consumption.

Table 1. Processor Architecture Comparison

The dedicated stack engine further reduces the number of micro-ops in flight by providing its own dedicated “adders”¹ to perform all stack updates. This approach frees up the integer execution units to work on other instructions in the pipeline. The result is fewer micro-ops per task, which reduces the power consumption per task.

An improved branch prediction algorithm significantly reduces overall power consumption. With fewer mispredicted branches, the number of wasted operations executed is reduced. Two additional predictors—the Indirect Branch Predictor and Loop Detector—also help to improve the prediction rate by tracking special program flows. Even the processor bus of the Pentium M helps to lower power consumption through circuit methods.

Performance Analysis

The benchmarks used by Dell Performance labs to compare the performance of the processors are broken up into four categories: office productivity, gaming, CPU, and workstation. The system configurations used in these benchmarks are shown in Table 2. The configurations were designed to isolate processor performance as much as possible by using identical hard drives and graphics cards.

	Pentium M System	Pentium 4 System
CPU	2.13-GHz Pentium M 533 FSB	3.8-GHz Pentium 4 800 FSB
Hard drive	Seagate 120-GB 7200 RPM	Seagate 120-GB 7200 RPM
Operating system	Microsoft® Windows® XP, Service Pack 2	Microsoft Windows XP, service pack 2
Video card-gaming	ATI X800XT 256 MB	ATI X800XT 256 MB
Video card-workstation	Nvidia FX 3400 256 MB	Nvidia FX 3400 256 MB
Memory	1 GB (2 x 512-MB DDR2 533 MHz)	1 GB (2 x 512-KB DDR2 533 MHz)

Table 2. Performance Benchmark System Configurations




Figure 1. Sysmark 2004 Relative Results

Office Productivity Benchmarks

The benchmarks analyzed were SYSmark 2004, Business Winstone 2004, Multimedia Content Creation Winstone 2004 and PCMark 2004. These results are summarized in Figures 1 and 2.²

SYSmark is a multithreaded application, so it scales well with physical or logical (via Hyper-Threading Technology) processors. This gives the Pentium 4, which supports Hyper-Threading Technology, an advantage over Pentium M, which does not. Between the subtests, the Pentium 4 shows the biggest lead on Internet Content Creation, largely due to its faster clock speed and a more-efficient floating point unit. The gains on the office productivity portion of the test can be linked to the faster frontside bus speed of the Pentium 4. With Hyper-Threading Technology disabled on the Pentium 4, the gap between the two processors shrinks. The Pentium 4 also receives a higher PC Mark CPU score due to its support for Hyper-Threading Technology. However, on Business Winstone 2004, the larger cache of the Pentium M proves to be a major benefit and helps it to move past the Pentium 4.

1. Adders are functional units in the microprocessor that perform arithmetic addition operations.

2. Actual performance will vary based on configuration, usage, and manufacturing variability.




Figure 2: Office Productivity Benchmark Relative Results

Gaming Benchmarks

All gaming benchmarks were run at a medium resolution of 1024x768 using identical desktop high-end graphics cards to isolate CPU performance. The use of a high-end desktop graphics card on a mobile system is not feasible outside the lab, given the limited thermal headroom available on mobile platforms. Similarly, the use of a high-end desktop Pentium 4 at full performance (that is, maximum frequency) is also not feasible on a mobile platform.

The gaming benchmark results are summarized in Figure 3. The Pentium M exhibits very good gaming performance compared to the Pentium 4, because of the larger cache and shorter pipeline of the Pentium M. (See comparison of Pentium M and Pentium 4 in Table 1.)²

In UT2004 and FarCry, the Pentium M wins on all the subtests. This is also the case in Doom3. In the multi-

threaded 3Dmark05, however, the Pentium 4 holds a huge lead over the Pentium M on the CPU subtest. The Pentium 4 support for Hyper-Threading Technology gives it a performance boost. This gap shrinks with Hyper-Threading Technology disabled. With the exception of the Aquamark and 3DMark05 CPU test, the other gaming benchmarks are not designed to take advantage of Hyper-Threading Technology and show very little performance gain with the feature enabled.

At the highest resolution tested (1600x1200), benchmarks are more graphics-intensive. At this resolution, performance is limited primarily by the graphics card. For this reason, there is little performance difference between the two processors, as shown in Figure 4.




Figure 4: Gaming Benchmark Relative Results at 1600x1200 Resolution

Overall, the frontside bus speed appears to be the limiting factor for the Pentium M and an increase in bus speed could potentially widen the gap in favor of the Pentium M.




Figure 3: Gaming Benchmark Relative Results

CPU Benchmarks

SPEC CPU2000 was the main benchmark used to compare performance between the two processors. The larger cache and higher instructions per clock rating of the Pentium M, coupled with its shorter pipeline, give it a slight advantage on the integer portion of this benchmark. The Pentium 4 processor performed better on floating-point benchmarks because of its more efficient floating-point unit and SSE performance, and higher clock speed. The SPECfp suite depends heavily on system bus bandwidth. As a result, the Pentium 4 has an inherent advantage because of its higher bus speed in a desktop configuration. The results are summarized in Figure 5.²




Figure 5: CPU Benchmark Relative Results

Workstation Benchmarks

The Pentium 4 processor had the performance lead on most of the benchmarks. 3D workstation applications are heavily dependent on bus latency and bandwidth. The lower bus speed of the Pentium M is a major reason for its lower performance. Even so, the Pentium M is reasonably competitive in these benchmarks be-

cause of its larger cache. The results are summarized in Figure 6.²

Battery Life and Weight

The performance-improving features built into the architecture of the Pentium M—micro-ops fusion, advanced branch prediction, and dedicated stack manager—also reduce overall power consumption. Shorter pipelines that do more work per clock require fewer transistors, and transistors that clock slower require less power. As a result, a mobile platform based on the Pentium M can be extremely power efficient.

In contrast, the longer pipelines and higher clock speeds of the Pentium 4 result in much higher power consumption and lower battery life. For this reason, portable computers based on the Pentium 4 processor must use higher-capacity batteries to achieve similar battery life, but this increases the weight of the system. Systems based on the Pentium M can be lighter because lower-capacity battery packs can be used without compromising battery life or performance. Also, due to lower power dissipation from the Pentium M, the thermal solutions required to keep the processor cool are lighter, which further reduces system weight.

Mobilemark 2002 was used to compare battery life between the Inspiron™ 5160 (Mobile Pentium 4 with HT) and the Inspiron 8600 (Pentium M). These systems have different chip sets, processors, display panels, and batteries. However, to get a feel for the battery life difference, we matched up other system components such as the graphics card, hard drive, and memory to keep the comparison as similar as possible. The results are summarized in Table 3. These results show that the




Figure 6: Workstation Benchmark Relative Results

	Inspiron 5160	Inspiron 8600
CPU	3.2-GHz Mobile Pentium 4, Hyper-Threading enabled	1.7-GHz Pentium M
Video Card	nVidia GeForce FX Go5200 32 MB	nVidia GeForce FX Go5200 32 MB
Memory	1 x 256 MB DDR 333-MHz	1 x 256 MB DDR 333-MHz
Panel	SXGA+	WUXGA
Battery	96 watt hours	72 watt hours
Maximum battery life ³	197 minutes	266 minutes
Average system power	29.24 watts	16.22 watts
Base weight	8.3 lbs	6.6 lbs

Table 3. Battery Life System Setup

Inspiron 8600 has better battery life and is lighter than the Inspiron 5160.³

Conclusion

This study shows that the Pentium M is competitive on almost all benchmarks. From a purely performance standpoint, its best showing is on the gaming benchmarks where it outperformed the Pentium 4 on several tests. In mobile platforms, where battery life and thermal issues take precedence, the Pentium M provides comparable performance with longer battery life and

lower weight. It also frees up power and thermal capacity, thus permitting the use of higher-performance components such as higher-end video cards. (See sidebar.) On most of the system-level and workstation benchmarks, the Pentium M is not far behind the desktop Pentium 4. With the increased use and popularity of portable computers, battery life and weight have become important deciding factors for most consumers. It is here that the Pentium M shines with its impressive battery life numbers and lower weight.

Each computer system has a fixed power and thermal limit or "budget." This budget is affected by the system's size and weight. System designers make various trade-offs to meet the power and thermal budget. The lower power requirements of the Pentium M processor frees up power and thermal capacity that can be used to provide features such as higher-end graphics or memory solutions that need the extra power, or a second hard drive for a RAID solution.

References

- http://www.intel.com/technology/itj/2003/volume07issue02/art03_pentiumm/p01_abstract.htm.

Steven Sexton, Systems Engineer with Dell System Performance Analysis, was an additional contributor to this white paper.

Disclaimer:

The performance tests and results in this document are measured using specific configurations within the Dell Performance labs. Actual performance results will vary based on the user's system configuration, applications, and manufacturing variability. These results should be used as a relative indicator only. Buyers should consider their own usage and consult with sources of performance data such as SPEC at www.spec.org, and BAPCo at www.bapco.com.

THIS WHITE PAPER IS FOR INFORMATIONAL PURPOSES ONLY, AND MAY CONTAIN TYPOGRAPHICAL ERRORS AND TECHNICAL INACCURACIES. THE CONTENT IS PROVIDED AS IS, WITHOUT EXPRESS OR IMPLIED WARRANTIES OF ANY KIND.

© 2005 Dell Inc. All rights reserved.

Trademarks used in this text: *Dell*, the *DELL* logo, and *Inspiron* are trademarks of Dell Inc.; *Intel* and *Pentium* are registered trademarks of Intel Corporation; *Microsoft* and *Windows* are registered trademarks of Microsoft Corporation. Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and names or their products. Dell Inc. disclaims any proprietary interest in trademarks and trade names other than its own.

3. Actual battery life will vary based on configuration, usage, and manufacturing variability.

https://web.archive.org/web/20070703180231/http://download.intel.com/research/platform/terascale/terascale_overview_paper.pdf



White Paper
Research at Intel

From a Few Cores to Many:

A Tera-scale Computing Research Overview

Editors:

Jim Held

Intel Fellow, Director, Intel Tera-scale
Computing Research Program

Jerry Bautista

Co-Director, Intel Tera-scale
Computing Research Program

Sean Koehl

Technology Strategist, Intel Tera-scale
Computing Research Program
Intel Corporation

1. Introduction

1.	Introduction	2
1.1	Motivation	3
1.2	New Opportunities	4
2.	Tera-Scale Computing Research Areas	5
2.1	Microprocessor Research Areas	5
2.1.1	Core Designs	5
2.1.2	Fixed-Function Units	6
2.1.3	Scalable On-Die Interconnect Fabric	6
2.1.4	Energy Management	6
2.2	Platform Research Areas	7
2.2.1	Cache and Memory Hierarchy	7
2.2.2	I/O	7
2.2.3	Virtualization and Partitioning	8
2.2.4	Execution Environment	8
2.3	Software Development Research Areas	8
2.3.1	Future Workloads	8
2.3.2	Programming Environment	10
2.3.3	New Technologies To Support Parallel Programming	10
3.	Summary	11

Contributors: Ali Adl-Tabatabai, Pradeep Dubey, Dave Dunning, Mike Espig, Ed Grochowski, Antonio Gonzalez, Scott Hahn, Ram Huggahalli, Jay Jayasimha, Akhilesh Kumar, Partha Kundu, Tim Mattson, Derek McAuley, Alberto Munoz, Chuck Narad, Don Newell, R. M. Ramanathan, Thom Sawicki, Sebastian Schoenberg, Ioannis Schoinas, John Shen, Jim Sutton, Manny Vara, Mona Vij.

Intel processors with two cores are here now, and quad-core processors are right around the corner. In the coming years, the number of cores on a chip will continue to grow, launching an era of vastly more powerful computers. These are the machines that will deliver teraflop performance with the efficient capabilities needed to handle tomorrow's emerging applications.

The Intel® Tera-scale Computing Research Program is Intel's overarching effort to shape the future of Intel processors and platforms. Intel researchers are already working on **over 100** R&D projects worldwide to address the hardware and software challenges of building and programming systems with dozens of energy-efficient cores with a sophisticated memory hierarchy. Currently, projects in this program span circuit technologies, microarchitecture, interconnects, memory, and software technologies.

Some may wonder why Intel is putting so much emphasis on tera-scale research. The reason is simple: tera-scale computing is a two-fold revolution, both in the capabilities that devices will have, and in the amount of innovation that will be required to handle tomorrow's advanced applications. The term itself—tera-scale—refers to the terabytes of data that must be handled by platforms capable of teraflops of computing performance. That's a thousand times more compute capability than is available in today's giga-scale devices.

Why such a leap forward? Because incremental improvements in performance and capabilities simply won't support real-time data mining across teraflops of data; artificial intelligence (AI) for smarter cars and appliances; virtual reality (VR) for modeling, visualization, physics simulation, and medical training; and other applications that are still on the edge of being science fiction. Also, data stores are becoming larger and more complex. In medicine, a full-body medical scan already contains terabytes of information. Even at home, people are generating large amounts of data, including hundreds of hours of video, thousands of documents, and tens of thousands of digital photos that need to be indexed and searched. Tera-scale computing is the way to bring the massive compute capabilities of supercomputers to everyday devices, from servers, to desktops, to laptops.

For example, with a tera-scale computer, you could create studio quality, photo-realistic 3-D graphics in real time. Or you could manage personal media better by automatically analyzing, tagging, and sorting snapshots and home videos. Advanced algorithms could be used to improve the quality of movies captured on older, low-resolution video cameras. An advanced digital health application might assess a patient's health by interpreting huge volumes of data in a scan and aid in making decisions in real time.

The essential aspect of tera-scale technologies—and the heart of Intel's research—is being able to do such complex calculations in real-time, primarily through the execution of multiple tasks in parallel. That is the fundamental requirement for the complex and compelling applications we will see in the future.

1.1 Motivation

In the last twenty years, Intel has delivered dramatic performance gains by increasing the frequency of its processors, from 5 MHz to more than 3 GHz, while at the same time, improving IPC (instructions per cycle). Recently, power-thermal issues—such as dissipating heat from increasingly densely packed transistors—have begun to limit the rate at which processor frequency can also be increased. Although frequency increases have been a design staple for the last 20 years, the next 20 years will require a new approach. Basically, industry needs to develop improved microarchitectures at a faster rate, and in coordination with each new silicon manufacturing process, from 45 nm, to 32 nm, and beyond.

For this new approach we can take advantage of Moore's law. Transistor feature size is expected to continue to be reduced at a rate similar to that in the past. For example, a 0.7x reduction in linear dimensions enables a 2.0x increase in the transistor density. Thus, we should

assume that with every process generation, we will be able to build chips with twice the number of transistors as on the previous process generation. New technologies, such as 3-D die stacking, may allow even greater increases in total transistor counts within a given footprint, beyond the increases made possible by improvements in lithography alone.




Figure 1. Scaling transistors. The number of transistors is expected to continue to double about every two years, in accordance with Moore's Law. Over time, the number of additional transistors will allow designers to increase the number of cores per chip.

With so many transistors available, we have already begun to design chips with multiple processor cores (also called CMP or chip-level multiprocessing). Instead of focusing solely on performing individual tasks faster, we will execute many more tasks in parallel at the same time. We will also distribute those tasks across a grouping of cores that work in a coordinated fashion. Three distinct trends motivate this shift:

- **Performance:** We can no longer simply increase the clock frequency (processor “speed”) at the same rate as we have in the past in order to increase performance. Power and thermal requirements are beginning to outstrip the benefits that faster clock frequencies offer. However, because the trajectory of Moore’s law will continue well into the next decade, we expect to continue doubling transistors every 18-24 months for the next several years. Parallel execution in multi-core designs will then allow us to take advantage of these greater transistor densities to provide greater performance.
- **Power consumption:** Many simple cores can be built within the same area as a small number of large complex cores. In addition, power consumption can be optimized by using multiple types of cores tuned to match the needs of different usage models. Also, cores that are not busy can be powered down to reduce power consumption during idle times. These advanced power-saving techniques are enabled by multiple cores working in a coordinated fashion.
- **Rapid design cycles:** Building tera-scale processors from standard, repeated tiles that are then integrated into a common infrastructure should allow more reuse of designs between generations of processors. This will also allow us to highly optimize the design of these tiles to further improve power utilization and performance.

Since Moore’s law is expected to continue to deliver more transistors every process generation, and since platform power and energy budgets will be increasingly limited, the trend is to deliver increased performance through parallel computing. Essentially, to achieve the desired improvements in performance without a corresponding increase in energy bills, we must increase the efficiency and number of cores on a chip, rather than increase clock frequency.

1.2 New Opportunities

Obviously, tera-scale architectures present many challenges. They also offer many unique opportunities through their highly integrated multi-core designs.

For example, consider on-chip core-to-core communication latencies and bandwidth. These will be orders of magnitude better than the chip-to-chip multiprocessing systems used today for parallel computing.

The number of cores in a chip can also reach levels well beyond those of traditional high-end servers. This can offer significant benefits in both performance and cost reductions over systems with traditional single or dual processor systems.

Finally, the very high computational power of a tera-scale processor, combined with the modularity of many cores, gives designers the flexibility to dedicate hardware resources (such as one or more cores) to specific functions. These could be system management or single-use devices, such as a Voice-over-IP appliance (for internet telephony), cryptography, or media encoding/decoding.

2. Tera-scale Computing Research Areas

Intel has identified several key attributes that will be required of future tera-scale platforms:

- **Programmability.** Without optimized software, tera-scale platforms will not live up to their potential. Platforms must effectively address the needs of new and existing programming models. This also includes software development/debug tools, as well as new performance benchmarks consistent with highly parallel execution.
- **Adaptability.** The platform must be able to change configuration to match varied usage and workloads, as well as adapt to changes in the hardware environment, such as from power and thermal factors.
- **Reliability.** The platform must preserve current levels of reliability or increase its reliability despite the increased complexity inherent in these platforms.
- **Trust.** The platform must provide a trustworthy environment, despite its flexibility and the complexity of its design.
- **Scalability.** The platform must deliver performance that increases in proportion to the number of cores, with hardware and software that also effectively scales.

To help develop platforms that meet these requirements, the Intel Tera-scale Computing Research Program encompasses projects around the globe, and addresses challenges in the three critical areas of microprocessors, platforms, and software development. Among others, these challenges include optimizing designs for highly parallel and multithreaded workloads, developing scalable on-die interconnects, redesigning cache and memory hierarchies, improving I/O, and identifying and exploiting parallelism via software development.

2.1 Microprocessor Research Areas

One of Intel's goals is to deliver a highly modular, scalable tera-scale processor architecture that can address the broad range of Intel platforms within future power constraints. There are four keys to delivering this architecture: optimized core designs, fixed-function units, scalable on-die interconnects, and effective energy management.

2.1.1 Core Designs

One of the differentiating aspects of tera-scale architectures is the use of a core microarchitecture that is optimized for highly-parallel, multithreaded workloads. Simultaneous multithreading can help deal with memory latencies and help reduce or eliminate out-of-order complexity.

At Intel, researchers are working to determine the optimal core design or designs needed for tera-scale architectures. Researchers are also working to identify the number of cores that can be effectively exploited in a given manufacturing process generation (such as 45 nm, 32 nm, or 22 nm). Of special importance is research aimed at including in each core special features/mechanisms that will facilitate the exploitation of thread-level parallelism.

Legacy software and inherently single-threaded algorithms present another challenge to tera-scale architectures. To address that issue, Intel researchers are exploring ways to incorporate heterogeneous general-purpose cores (both single-thread and multithread optimized cores) into tera-scale architectures. In addition, researchers are exploring the use of ensembles of simple cores to accelerate single threads through thread-level speculation and other advanced techniques.

By varying the mix of functional elements, the architecture can allow designers to create multiple implementations that match specific market needs. If the processor uses regular tiles that consist of one or more cores, cache, a router, and other supporting hardware, different aspects of a layout could be reused from processor to processor.




Figure 2. Future tera-scale chips will use an array of tens to hundreds of cores with reconfigurable caches, as well as special-purpose hardware accelerators.

2.1.2 Fixed-Function Units

The demands of power efficiency suggest that in some instances it will be worthwhile to create fixed-function units. Fixed-function units are processing elements dedicated to a particular workload. They will perform better for that task than a general-purpose execution unit would perform (for example, a variable-length bit manipulation or video decode logic). These functions may be fixed logic or programmable engines that are optimized to a particular purpose, such as graphics engines to accelerate texture processing. Other possibilities under research include various network accelerators, cryptography engines, and physics engines.

With fixed-function units, a given tera-scale platform can have a processor with a mix of element types matched to the significant workloads in the segment it is to serve. Both the number of each and the relative number of each type could be varied.

2.1.3 Scalable On-Die Interconnect Fabric

A tera-scale processor needs to connect, not only to a large number of cores, but also to caches and special-purpose hardware, such as graphics units. Basically, as more elements are packed onto a tera-scale chip, there is a correspondingly greater need for each of these units to communicate with each other. Studies have shown that the interconnect design and microarchitecture will play a significant role in determining the net performance of a tera-scale architecture.

In particular, tera-scale architecture will require a modular on-die (or on-chip) interconnect “fabric” that is scalable to support a variable number of cores. Interconnects must also be robust in the face of failures, and have an architectural lifetime that can span several generations of processor designs. Finally, the interconnects must have efficient power management that scales down power consumption to match the more efficient utilization of the cores. On-chip interconnect networks are expected to have several advantages over today’s off-chip core interconnects:

- On-chip wires are cheap and plentiful compared to off-chip networks in printed circuit boards.
- On-chip interconnects can often be routed on top of some types of circuit structures (such as caches), so that they consume little on-die real-estate.
- Given the smaller distances on-die interconnects must span, such interconnects are generally more responsive and power efficient than off-die interconnects.
- On-die bus widths can be much wider than those on printed circuit boards, allowing for more efficient, lower modulation speeds. In other words, bandwidths can be increased by widening the bus and actually slowing the modulation speed to save power.

Because on-die interconnects can support a wide number of cores, they can serve many market segments, ranging from large server devices to small mini-notebook processors. The interconnect architecture would support modular designs (with a standard interface to the various compute elements), and offer performance headroom to scale over a decade in order to support several process generations. Such characteristics are typical of platform interconnects such as PCI, which further highlights the essential nature of a tera-scale processor as a platform/system on a chip.

2.1.4 Energy Management

A tera-scale architecture would offer new opportunities for energy management. For example, the availability of numerous, modular cores will allow the system to migrate work as needed to balance out utilization across the tera-scale processor, or to match workload to the type of core. Effective management of thermal loads in individual components of a tera-scale processor will also improve the overall reliability, as well as improve overall performance by evenly spreading out the thermal load.

At the circuit level, Intel is researching many critical areas. These include designs for circuits for on-die caches and on-die memory that are much more efficient in their use of power. Prototype register file circuits are already showing up to a 3X improvement in the speed of memory accesses, and up to 4X to 5X reduction over today’s circuits in the power they consume. Intel’s goal is a 10X improvement in performance per watt over the next 10 years.

Tera-scale processors may also have the ability to throttle the energy efficiency for a given compute task. This means noncritical or highly parallel workloads could be run at a reduced frequency, and thus much lower energy consumption. For critical code or segments of highly serial code, performance could be boosted by temporarily throttling up the frequency.

2.2 Platform Research Areas

Balancing a system with the tremendous compute density of a future tera-scale processor will be challenging. For example, getting data in and out of a tera-scale processor in an efficient manner requires balancing cache and memory bandwidth, I/O bandwidth, and connectivity. Essentially, a tera-scale processor needs to be able to move tera-bytes of data onto and then later off the chip. As the number of cores grows, it will be increasingly challenging to provide I/O with a performance level that scales to the number of cores.

Tera-scale platform hardware will also need to support an execution environment which provides power management, partitioning, resilience, and other features of the full system.

2.2.1 Cache and Memory Hierarchy

One major set of changes to platform design will be in the memory hierarchy. Specifically, the cache hierarchy for a tera-scale device will include private first-level caches (L1 caches) per core, as well as several levels of shared cache (L2 cache). Research in these areas includes work on shared distributed caches, cache policies (including data-specific policies), and cache partitioning.

In addition, a tera-scale processor will most likely require an additional level in the memory hierarchy to match its bandwidth requirements to the system memory characteristics. Here, researchers are exploring the use of 3-D stacking to provide a large, low-latency, last-level cache using SRAM or DRAM.

2.2.2 I/O

Tera-scale platforms also create new challenges associated with I/O. For example, the I/O of a tera-scale device will have to scale to match the compute density (the number of transistors for a given footprint) of its architecture. This means the device must be able to handle the potential aggregate bandwidth required when all cores are operating at the same time.

While providing the potential aggregate bandwidth raises considerable challenges, issues of connectivity and sharing will be even more difficult to resolve. This is because the level of I/O sharing is exaggerated due to the presence of more cores and threads per processor.

At the physical layer, copper I/O interconnects must scale to support hundreds of gigabits per second. This performance is required to support the expected tera-scale workloads. Intel is already researching ways to scale existing copper I/O interconnects to tens of gigabits per second per line.

Hardware and software mechanisms for device replication and sharing will also need to scale. These mechanisms will have to support dozens to hundreds of threads or virtual machines that are concurrently sharing a device. The challenge is that most practical configurations incorporating tera-scale processors will be limited to one or a small number of disks. To address this issue, Intel researchers are also looking at ways to incorporate into storage subsystems the hardware and software mechanisms that would provide the necessary bandwidth and I/O per second.

A tera-scale I/O architecture must also be flexible and versatile enough to balance performance, cost, and power across the many workloads, benchmarks, and usage models. Balanced I/O refers to providing I/O services in proportion to the compute demands of applications running on the platform, so that the I/O subsystem does not critically limit the application's performance. To be balanced, I/O as a subsystem must also meet the above requirement with a minimal cost in silicon, software, and energy consumption.

Intel has been conducting extensive research on ways to provide performance I/O for tera-scale devices through technologies such as CMOS radio and silicon photonics. CMOS radios would provide the capability to integrate flexible, multistandard wireless capabilities into components such as chipsets. Likewise, Intel's recent advancements in silicon photonics are aimed at making high-bandwidth fiber-optic links affordable for connections in and around desktops and servers.

2.2.3 Virtualization and Partitioning

Intel's vision for tera-scale platforms includes support for both hardware partitioning and virtualization. These capabilities will support special usage models, as well as improve robustness and trust through tamper-resistant, hardware-based isolation. The virtualized platforms will have full support for efficient access to such partitions from and to I/O devices.

Applications that process real-time media may require performance guarantees in order to provide the required quality of service (QoS) to the user. Because of this, performance isolation among partitions or among virtual machines will be important, especially for client tera-scale platforms.

2.2.4 Execution Environment

The execution environment encompasses the system software and firmware that controls access to shared resources on the platform. These resources include compute resources (cores), storage resources (file systems and disk), and memory.

The software and firmware elements of the execution environment include the operating system, hypervisor or virtual machine monitor, and any firmware that is used to partition the system. As the tera-scale architecture increases the parallel nature of workloads, the way elements access shared system resources must change to ensure the availability of the resources. Each element is responsible for making sure that its workload receives access in such a way that meets the needs of other workloads running on tera-scale platforms. In other words, each element must share resources appropriately in order to allow other components to access the same resources.

As with other aspects of a tera-scale architecture, the elements of the execution environment must scale with the number of cores. In addition, the elements must be able to adapt to platforms where resources are being dynamically added or removed.

2.3 Software Development Research Areas

Developing software to take advantage of tera-scale hardware includes two of the greatest challenges for tera-scale computing.

The first challenge is to ensure that there are compelling applications and workloads that exploit the massive compute density, using parallel algorithms and potentially masses of data. Intel researchers are already working to identify and enable such workloads and algorithms. The goal here is to distill the characteristics of potential workloads and algorithms, then develop technologies that facilitate application development and scale them onto tera-scale architectures platforms.

The second challenge is that multiprocessing adds a time dimension that is extremely difficult for software developers to cope with and extremely difficult for validation engineers to test. The large-scale concurrency of such devices will only increase software development challenges—such as implementing correct synchronization among threads without deadlock and race conditions in accessing shared memory. These are very subtle errors that are hard to reproduce. Intel is researching software technologies such as transactional memory to address these types of programming needs. In addition, Intel is working with universities worldwide to expand curricula to include multithreaded software development and train future generations of developers.

Certainly not all legacy applications will be rewritten to take advantage of a parallel programming environment. To address the legacy issue, Intel researchers are working on compiling techniques and runtime environments that automatically parallelize existing, single-threaded code. Such techniques have limited benefits relative to building the code from the ground up with a parallel execution model in mind. However, the techniques may allow single threaded applications to take some advantage of the higher performance capabilities of multi-core, tera-scale devices.

2.3.1 Future Workloads

Future usage models will define the requirements of tera-scale platforms. As such, industry needs effective models for what the software workloads will be and how those workloads will fit together. Intel researchers are currently working to identify and model the emerging workloads and how they will interact with specific features of the tera-scale architecture.

While existing single-threaded workloads are important, new applications will be enabled by tera-scale processors. Compute densities on future tera-scale processors will be significantly greater than the densities found on current microprocessors. This will help developers shift applications that were once possible only on high-end supercomputers, onto laptops or even handheld devices.

To this end, Intel is investigating classes of processing capabilities that will be needed to handle tera-scale workloads. Intel classifies these processing capabilities into three fundamental types: recognition, mining and synthesis (RMS).

- **Recognition:** Machine-learning capabilities that allow computers to examine data and images and construct mathematical models based on what they identify. An example might be a model for the face of a specific person.
- **Mining:** The capability to sift through large amounts of real-world data related to the patterns or models of interest. Put more simply, it is the ability to find an instance of a specific model amidst a large volume of data. For example, mining could mean finding a particular person's face from a large number of images of various resolutions, lighting environments, and so on.
- **Synthesis:** The capability to explore theoretical scenarios by constructing new instances of a model. For example, this could be projecting what the target person's face might look like if they were younger or older.

An example of an RMS workload is computer-assisted care of the infirm or the elderly. With standard web cameras wirelessly connected and distributed around a house, a tera-scale computer could recognize the person under supervision and model expected behavior. For example, the computer could help make sure the person self-medicates at appropriate intervals, gets a glass of water to drink after so much time has passed, or doesn't leave the stove on and unattended for more than a few minutes. Such modeling applications could help reduce health risks for the elderly and allow for increased self-reliance. The challenge with such applications is that the compute requirements

are intense. Advanced image processing, inferential logic, anticipation of potentially dangerous scenarios, and the security to ensure privacy for sensitive information—these require extremely intensive compute performance.

RMS workloads are of particular interest to Intel since they show a great deal of promise to deliver user value via a common set of algorithms that scale. The RMS application model appears to be general in nature, applying to a very large number of compute needs that span high-performance computing, digital content creation, computer vision, and artificial intelligence. Essentially, these workloads have certain system and architectural commonalities. By deeply studying the underlying RMS algorithms and kernels, Intel researchers can deduce the key architectural requirements that benefit this class of applications. In turn, this will help researchers turn an understanding of RMS into improved platform architectures for the future.

Regardless of how computing evolves over the next decade, people will still use web servers, office productivity applications, and databases. Intel continues to study these more traditional workloads to understand how they relate to future tera-scale architectures.




Figure 3. Designing realistic prototype applications allows Intel researchers to test multithreaded scalability and feed key learning into future hardware designs.

2.3.2 Programming Environment

One of the more difficult challenges may be efficiently writing software that gains the benefit of multi-core processing. In the tera-scale future, software should be designed to use available parallelism to gain the performance benefit of the increased numbers of cores. This requires that software developers design parallel programs, a traditionally time-consuming and error-prone task which requires developers to think differently than the way they do today. Teaching mainstream and future developers to identify and then effectively exploit parallelism is something Intel must foster if these skills are to move from a narrow domain of high-performance computing (HPC) experts into the mainstream.

One of the primary issues is that almost all of the elements associated with the programming environment will have to change in order to accommodate scalability and concurrency. Compilers and tools such as debuggers and performance analyzers will need to support new concurrency constructs. The programming environment itself will need to provide new language constructs, as well as new tools to help the programmer uncover parallelism and avoid parallel programming bugs (such as deadlocks and race conditions). Language runtimes must also scale, provide new concurrency primitives

for the compiler, and expose new interfaces to tools. Finally, libraries will have to provide new APIs that make it easier to develop robust and scalable applications. Beyond the more common shared-memory model of SMP (symmetric microprocessor) programming, Intel is investigating tools to support programming models that have been successfully employed in existing parallel programming environments. These environments include the streaming programming model used with GPUs (graphics processing units), and the distributed-processing model based on message processing that is common in HPC.

2.3.3 New Technologies to Support Parallel Programming

To simplify software development, Intel is researching new features for tera-scale devices. For example, transactional memory simplifies parallel programming by reducing the need for software developers to manage explicit locks. A key focus of Intel's research is finding other hardware features that will simplify parallel programming.

One of the features of tera-scale architectures will be dedicated partitions that appear as devices to regular software. These partitions will provide functions such as system management and I/O acceleration (such as network protocol processing). The architectures may also include hardware support for lightweight message passing for a distributed-computation model that is familiar to HPC software developers. This type of streaming has proven to be an effective programming model for workloads such as graphics and media processing. Tera-scale architectures will exploit that familiarity by providing hardware support. For example, the architectures will support the cache behavior desired, as well as coprocessor instruction access to fixed-function media acceleration.

In addition, as with the MMX extensions for multimedia applications, there will be opportunities to extend the ISA (Instruction Set Architecture) to better support emerging RMS workloads.


3. Summary

Intel foresees future devices revolving around tera-scale architectures that use tens to hundreds of cores to process massive amounts of information in parallel. By making these devices scalable, adaptable, and programmable, Intel can help lead industry into a new era where more immersive, interactive, and useful applications become possible.

In a tera-scale world, there will be new processing capabilities for mining and interpreting the world's growing mountain of data, and for doing so with even greater efficiency. Intelligent agents could advise users in real-time on stock trades and other financial decisions. Such agents could search massive collections of digital videos to find specific people or events, and even edit a new video based on what the user wants to see. For gamers, there is the obvious benefit of photo-realistic, real-time graphics. However, even those benefits aren't just for gamers anymore. Interactive virtual environments are now being developed for both collaboration and education, such as learning a language by interacting with virtual native speakers, or training a doctor to deal with an emergency situation on a simulated human body.

There are people who still question a multi-core, tera-scale future, and why tomorrow's applications need so many threads. The answer is that those advanced, intelligent applications require supercomputing capabilities, and the accompanying parallelism that allows those applications to be processed in real-time. This is a massive shift in what mainstream devices—servers, desktops, and mobile—can do, and it requires an equally massive shift in hardware and software.

Tera-scale computing is not a new goal, but the fulfillment of Platform 2015. Intel's long-range vision for the collective evolution of computational technologies, interfaces, and infrastructures, and the architectural innovation and core competencies will enable that evolution. Through the Tera-scale Computing Research Program, Intel has dedicated a significant amount of effort towards developing the technologies to realize this vision. The resulting discoveries and successes will not only shape the future of Intel micro-architecture, but will guide the capabilities of the underlying platforms, and allow the possibilities of the future to become reality through revolutionary applications.



www.intel.com

References

Xiao-Feng, Li Zhao Hui Du, Chen Yang, Chu-Cheow Lim, Tin-Fook Ngai, "Speculative parallel threading architecture and compilation", pp 285- 294, International Conference Workshops on Parallel Processing (ICPP 2005), 2005.

Murali Annavaram, Ed Grochowski, John Shen, "Mitigating Amdahl's Law through EPI Throttling," pp. 298-309, 32nd Annual International Symposium on Computer Architecture (ISCA'05), 2005.

D. Nelson, C. Webb, D. McCauley, K. Raol, J. R. II, J. DeVale, and B. Black, "A 3D Interconnect Methodology Applied to iA32-class Architectures for Performance Improvements through RC Mitigation," in Proceedings of the 21st International VLSI Multilevel Interconnection Conference, Waikoloa Beach, HI, USA, September 2004.

Paul Reed, Gus Yeung, and Bryan Black. "Design Aspects of a Microprocessor Data Cache using 3D Die Interconnect Technology." In Proceedings of the International Conference on Integrated Circuit Design and Technology, pages 15-18, Austin, TX, USA, May 2005.

L. Hsu, R Iyer, S Makineni, S Reinhardt, D Newell, "Exploring the cache design space for large scale CMPs", pp 24-33, ACM SIGARCH Computer Architecture News, Volume 33, Issue 4 (November 2005).

Berna L. Massingill, Timothy G. Mattson, and Beverly A. Sanders; "Reengineering for Parallelism: An Entry Point into PLPP (Pattern Language for Parallel Programming) for Legacy Applications"; Proceedings of the Twelfth Pattern Languages of Programs Workshop (PLoP 2005), 2005.

Pradeep Dubey, "Recognition, Mining and Synthesis Moves Computers to the Era of Tera" in Intel Technology Journal, February, 2005.

Ali-Reza Adl-Tabatabai, Brian T. Lewis, Vijay Menon, Brian R. Murphy, Bratin Saha, Tatiana Shpeisman, "Compiler and runtime support for efficient software transactional memory" pp 26-37, Proceedings of the 2006 ACM SIGPLAN conference on Programming language design and implementation.

Y. Chen, Q. Diao, C. Dulong, W. Hu, C. Lai, E. Li, W. Li, T. Wang, and Y. Zhang. Performance scalability of data mining workloads in bioinformatics. Intel Technology Journal, 09(12):131-142, May 2005.

More Info

Learn more about tera-scale computing research. Visit www.intel.com/go/terascale.

Learn more about other Intel research. Visit www.intel.com/research.

Copyright © 2006 Intel Corporation. All rights reserved. Intel, Intel logo, Intel. Leap ahead., and Intel. Leap ahead. logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

*Other names and brands may be claimed as the property of others.
Printed in the United States. 0906/MLG/HBD/PDF 315297-001US

https://web.archive.org/web/20091007000925/http://download.intel.com/pressroom/kits/events/idffall_2009/pdfs/22nm_factsheet.pdf



News Fact Sheet

CONTACT: Megan Langer
503-712-4305
megan.e.langer@intel.com

Intel Developer Forum 22nm News Facts


Sept. 22, 2009 — Intel Corporation is holding its Intel Developer Forum in San Francisco from Sept. 22-24. Paul Otellini, Intel president and CEO, today displayed the world's first working chips built on 22nm process technology. Intel continues to relentlessly pursue Moore's Law and provide its benefits to end users. Marking the third generation of high-k metal gate transistors, this event comes 2 years after Intel demonstrated working test circuits on the previous 32nm generation, and validates that Moore's Law continues well past the point when experts had predicted that the industry would hit a scaling wall.

- SRAMs are used as test vehicles to demonstrate technology performance, process yield and chip reliability prior to ramping processors and other logic chips that will use the given manufacturing process.
- Intel is now in full development mode on 22nm and on pace to continue the company's "tick- tock" model into the next generation.
- The 22nm test circuits include both SRAM memory and logic circuits to be used on 22nm microprocessors.
- SRAM cells of 0.108 and 0.092 square microns function in an array totaling 364 million bits. The 0.108 square micron cell is optimized for low voltage operation. The .092 square micron cell is optimized for high density and is the smallest SRAM cell in working circuits reported to date.
- The test chip packs 2.9 billion transistors, at approximately double the density of the previous 32nm generation, in an area as small as a fingernail.
- The 22nm dimensions are patterned with exposure tools using light with a wavelength of 193nm, a remarkable testament to the ingenuity of Intel's lithography engineers.
- This 22nm technology continues to deliver the promise of Moore's Law: smaller transistors, improved performance/watt and lower cost per transistor.

Intel and the Intel logo are trademarks of Intel Corporation in the United States and other countries.

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





The Programmable Solutions Company®

Innovate, differentiate, and win

At Altera, we've been promoting innovation—yours and ours—ever since we invented the world's first programmable logic device (PLD) 25 years ago. Today, we're excited to begin introducing 40-nm devices as part of our custom logic portfolio.

Our new 40-nm FPGAs and ASICs will give you the freedom and flexibility to bring your design concepts to life, without having to sacrifice one capability for another. Collaborating closely with our manufacturing partner, Taiwan Semiconductor Manufacturing Company (TSMC), we're building 40-nm devices that address a range of concerns, from power consumption to performance to cost.

Serving a variety of industries—including automotive, broadcast, computer and storage, consumer, industrial, medical, military, test and measurement, wireless, and wireline—our comprehensive portfolio includes:

- The industry's most advanced FPGA, CPLD, and ASIC technologies
- A fully integrated software development tool
- Versatile embedded processors
- Optimized intellectual property (IP) cores
- A variety of development kits
- Reference designs and design examples

Turn the page to learn more about Altera® custom logic solutions. For white papers, webcasts, application notes, and more resources, visit us at www.altera.com.

Contents

• MAX CPLD series	4
• Cyclone low-cost FPGA series	7
• Arria GX FPGA family	13
• Stratix high-end FPGA series	15
• HardCopy ASIC series	29
• Nios II embedded processors	38
• Intellectual property solutions	41
• Quartus II design software	46
• Development kits	49
• Configuration devices	57
• Training	61
• Package dimensions	64
• Ordering codes	67
• Product websites and information	69

MAX CPLD series

Altera's market-leading MAX® series of CPLDs are world-class, low-cost devices designed for virtually any digital control function. As non-volatile, single-chip solutions, MAX CPLDs are easy to incorporate into your system. With the devices, you can solve board-level issues such as insufficient I/O pins on a processor, manage the power-up sequence or configuration of other more complex devices, or inexpensively convert incompatible interfaces (a.k.a. "glue logic"). Designed to be hassle-free with intuitive device behavior and software, MAX CPLDs give you the freedom to focus on your more complex design challenges.

Key features


- Low cost
- Zero-power options
- Ultra-small packages
- Instant-on and non-volatile
- In-system programmability (ISP)
- Free Quartus® II Web Edition software support
- Free ModelSim®-Altera® Web Edition software support

MAX II

The market's lowest-cost CPLDs

MAX II devices are based on a groundbreaking architecture that combines the best of FPGAs and CPLDs. The performance and density benefits of a four-input look-up table (LUT) architecture are merged with the usability and cost advantages of a non-volatile architecture. The result? A first-in-its-class, new architecture that sets standards for cost, power consumption, performance, and density for CPLDs.

MAX II application areas



With MAX II CPLDs, you can reduce your system costs by integrating a larger amount of logic into a single device. This instant-on, non-volatile device family targets general-purpose logic applications, and comes in three variants: MAX II, MAX IIG, and MAX IIZ CPLDs.

The zero-power MAX IIZ CPLDs are the newest addition to the family. Ideal for portable and other power-, space-, and price-constrained applications, MAX IIZ devices come in ultra-small packages packed with more logic and I/O resources than are offered by traditional macrocell-based CPLDs in the same package size.

MAX II family features summary

Cost-optimized architecture	Revolutionary MAX II architecture delivers 4X the density at half the price of competing CPLDs.
Low power/zero power	Reduces power consumption and increases system reliability. The new zero-power MAX IIZ devices deliver the industry's lowest power consumption (standby and dynamic).
Highest-density CPLDs	Implements more applications in a single, low-cost device.
Small packages	Integrates on average 50 percent more user I/O pins and logic per board area (mm ²) than competing CPLDs with 0.5-mm pitch ball-grid array (BGA) packages. MAX IIZ devices come in two new ultra-small Micro Fineline BGA (MBGA) packages that deliver up to 6X the density and 3X the I/O resources in the same space as traditional macrocell-based CPLDs.
Non-volatile and instant-on functionality	Reduces system cost and board space with a single-chip solution.
User flash memory	Minimizes system cost and chip count by integrating discrete serial or parallel non-volatile storage onto MAX II devices.
Real-time ISP	Reduces maintenance costs by enabling updates while the device is in operation.
MultiVolt core	Operates with a 1.8V, 2.5V, or 3.3V power supply, minimizing power rails and simplifying board design.
MultiVolt I/O interface	Interfaces seamlessly to other devices at 1.5V, 1.8V, 2.5V, or 3.3V logic levels.
Parallel flash loader	Simplifies board management by using MAX II devices to configure external non-JTAG-compliant flash devices.

General-purpose CPLDs

Altera's 3.3V MAX 3000A devices are cost-optimized for high-volume applications, while the 5.0V, 3.3V, and 2.5V MAX 7000 families offer world-class, high-performance solutions for a broad array of applications.

The non-volatile, EEPROM-based MAX 3000A and MAX 7000 families provide instant-on capability and offer densities from 32 to 512 macro-cells. These devices support ISP and can be easily reconfigured in the field.

MAX 3000A family features summary

Low price per macrocell	Ideal for low-cost, high-volume applications.
4.5-ns propagation delays	Provides fast system performance.
5.0V tolerant I/O pins	Inherently interfaces to 5.0V devices.
Commercial and industrial temperature	Reduces overall system cost for temperature-sensitive applications.

MAX 7000 family features summary

4.5-ns propagation delays	Provides fast system performance.
5.0V tolerant I/O pins	Inherently interfaces to 5.0V devices.
Support for advanced I/O standards	Supports GTL+ and SSTL-2/-3 I/O standards (MAX 7000B CPLDs).
Programmable power-saving mode	Reduces power consumption by over 50 percent.
Commercial, industrial, and extended temperatures	Provides support for all environmental conditions.

MAX CPLD series package and I/O matrix

- Device available in commercial temperature. Contact Altera for industrial temperature.
- Device available in commercial and industrial temperatures.
- Device available in commercial and industrial, and qualified to extended temperatures.
- 36 Number indicates available user I/O pins.
- Vertical migration (Same V_{CC}, GND, ISP, and input pins).

		MAX II CPLDs 3.3V, 2.5V, 1.8V				MAX 3000A CPLDs 3.3V					MAX 7000AE CPLDs 3.3V				
		EPM240/GIZ	EPM570/GIZ	EPM1270/G	EPM2210/G	EPM3032A	EPM3064A	EPM3128A	EPM3256A	EPM3512A	EPM7032AE	EPM7064AE	EPM7128AE	EPM7256AE	EPM7512AE
Density and speed	Macrocells ¹	192	440	980	1,700	32	64	128	256	512	32	64	128	256	512
	Logic elements (LEs)	240	570	1,270	2,210	—	—	—	—	—	—	—	—	—	—
	Pin-to-pin delay (ns) ²	4.7, 4.7, 7.5	5.4, 5.4, 9.0	6.2, 6.2	7.0, 7.0	4.5, 7.5, 10	4.5, 7.5, 10	5.0, 7.5, 10	7.5, 10	7.5, 10	4.5, 7.5, 10	4.5, 7.5, 10	5.0, 7.5, 10	5.5, 7.5, 10	7.5, 10, 12
PLCC (L) ³	44 pin					34	34				36	36			
	84 pin												68		
TQFP (T) ⁴	44 pin					34	34				36	36			
	100 pin	80	76				66	80				68	84	84	
	144 pin		116	116				96	116				100	120	120
PQFP (Q or R) ⁵	208 pin								158	172				164	176
BGA (B) ⁶	256 pin														212
FBGA (F) ⁷	100 pin	80	76									68	84	84	
	256 pin		160	212	204			98	161	208			100	164	212
	324 pin				272										
MBGA (M) ^{8,9}	68 pin	54													
	100 pin	80	76												
	144 pin		116												
	256 pin		160	212											

¹ Typical equivalent macrocells for MAX II devices

² MAX IIZ data is preliminary

³ Plastic J-lead chip carrier

⁴ Thin quad flat pack

⁵ Plastic quad flat pack

⁶ Ball-grid array (1.27 mm)

⁷ FineLine BGA (1.0 mm)

⁸ Micro FineLine BGA (0.5 mm)

⁹ MAX IIZ devices are only available in commercial temperatures

MAX CPLD series features

		MAX II CPLDs 3.3V, 2.5V, 1.8V				MAX 3000A CPLDs 3.3V					MAX 7000AE CPLDs 3.3V				
		EPM240/G/Z	EPM570/G/Z	EPM1270/G	EPM2210/G	EPM3032A	EPM3064A	EPM3128A	EPM3256A	EPM3512A	EPM7032AE	EPM7064AE	EPM7128AE	EPM7256AE	EPM7512AE
Features	User flash memory (Kbit)	8				-					-				
	Boundary scan JTAG	✓				✓					✓				
	JTAG ISP	✓				✓					✓				
	Fast input registers	✓				-					✓				
	Programmable register power-up	✓				-					✓				
	Programmable ground pins	✓				-					✓				
	Open-drain outputs	✓				✓					✓				
	Programmable pull-up resistors	✓				-					-				
	Bus hold	✓				-					-				
	JTAG translator	✓				-					-				
	Real-time ISP	✓				-					-				
	0.5-mm BGA packages ¹	✓				-					-				
Core voltage and I/O options	Core voltage (V)	1.8				3.3					3.3				
	MultiVolt core (V)	3.3, 2.5, 1.8 ²				-					-				
	MultiVolt I/Os (V)	3.3, 2.5, 1.8, 1.5				5.0, 3.3, 2.5					5.0, 3.3, 2.5				
	I/O power banks	2	2	4	4	1					1				
	Maximum I/O pins	80	160	212	272	34	66	98	161	208	36	68	100	164	212
	Maximum output enables	80	160	212	272	6	6	6	6	10	6	6	6	6	10
	Transistor-to-transistor logic (TTL) (5.0V tolerance)	-	-	✓ ³	✓ ³	✓					✓				
	LVTTL/LVCMOS	✓				✓					✓				
	32-bit, 66-MHz PCI compliant	-	-	✓	✓	-					-				
	GTL+/SSTL-2, SSTL-3, all class I and class II	-				-					-				
	Schmitt triggers	✓				-					-				
	Programmable slew rate	✓				✓					✓				
	Programmable drive strength	✓				-					-				

¹ Package not available for EPM2210/G device

² MAX IIG and MAX IIZ devices require 1.8V core voltage supply

³ An external series resistor must be used for 5.0V tolerance

FOR MORE INFORMATION

MAX II CPLDs www.altera.com/max2
 MAX 3000A CPLDs www.altera.com/products/devices/max3k
 MAX 7000 CPLDs www.altera.com/products/devices/max7k
 Training www.altera.com/training

Cyclone low-cost FPGA series

For your cost-sensitive, high-volume applications, Altera offers our Cyclone® FPGA series—the industry’s only FPGAs designed from the ground up for low cost. Each family member is individually optimized for cost, and delivers a high-volume solution that’s competitive with ASICs and ASSPs. This series of FPGAs—including the 65-nm Cyclone III, 90-nm Cyclone II, and 130-nm Cyclone families—delivers a customer-defined feature set, industry-leading performance, and the lowest power consumption.


Key features

- The industry’s lowest-cost FPGAs
- High-performance digital signal processing (DSP)
- Low-cost embedded processing
- Free Quartus® II Web Edition software support
- Free ModelSim®-Altera® Web Edition software support
- Available in commercial, industrial, and automotive temperature grades

Cyclone III 65-nm low-cost FPGAs

Cyclone III FPGAs, the newest offering in this series of low-cost devices, features an unprecedented combination of low power, high functionality, and low cost to deliver more, sooner, and for less—even for your most cost-sensitive, high-volume applications. Built on TSMC’s 65-nm low power (LP) process technology, Cyclone III FPGAs were designed to provide customers with the flexibility and application-optimized features to enable the highest levels of design possibilities and productivity while meeting the most stringent cost and power budgets. What’s more, this can be accomplished without the high NRE costs associated with ASICs.

Cyclone III floorplan



Because our engineers defined, designed, and developed Cyclone III FPGAs with customer feedback in mind, the devices provide the ideal features to meet your needs. The logic-, memory-, and DSP-rich architecture enables you to enhance system integration for your applications. From video and image processing to displays and wireless applications, the opportunities for Cyclone III devices are unlimited.

Cyclone III family features summary

Cost-optimized architecture	Offers from 5,136 to 119,088 LEs—70 percent more than the Cyclone II FPGA family for enhanced system integration.
Low-power architecture	Altera’s innovative power-saving features and TSMC’s 65-nm LP process technology come together to minimize both standby and dynamic power consumption. Quartus II design software provides a power-aware design flow to enable optimization for minimal power usage. Compared to Cyclone II FPGAs, Cyclone III devices deliver up to 50 percent lower power.
Embedded memory	Up to 4 Mbits of embedded memory for memory-intensive applications such as video line buffers.
Embedded multipliers	Up to 288 dedicated 18-bit x 18-bit multipliers for high-bandwidth parallel processing applications.
External memory interfaces	Support for high-speed external memory interfaces including DDR, DDR2, SDR SDRAM, and QDR II SRAM at up to 400 megabits per second (Mbps). The autocalibrating external memory interface PHY feature eases timing closure and eliminates variations over process, voltage, and temperature (PVT).
Robust clock management	Up to 4 phase locked loops (PLLs) per device with 5 outputs per PLL providing up to 20 global clocks. Supports dynamic reconfiguration for frequency and phase changes.
Differential signaling	Supports up to 875-Mbps receive and 840-Mbps transmit LVDS signaling. Ability to use reduced swing differential signaling (RSDS), LVDS, and point-to-point differential signaling (PPDS) without external resistors.
Commodity parallel configuration	Support for low-cost configuration options.
Remote system upgrade	Allows storage of multiple configuration images in a single configuration device for field upgrades. Automatic error correction circuitry restores factory image if error detected.

Cyclone III family features summary (continued)

Automatic single event upset (SEU) detection circuitry	Features automatic SEU detection circuitry utilizing 32-bit cyclic redundancy check (CRC).
Hot socketing and power sequencing	Robust on-chip hot-socketing and power-sequencing support that ensures proper device operation independent of the power-up sequence.
MegaCore® IP library	Shorten design time using a broad portfolio of more than 200 Altera and partner intellectual property (IP) cores. Altera IP can be evaluated in hardware before purchase.
Nios® II embedded processor support	Nios II embedded processors offer an ideal replacement for low-cost discrete microprocessors to reduce cost and increase flexibility. Nios II soft processor support for Cyclone III FPGAs offers over 100-DMIPS performance.
Free software support	The free, downloadable Quartus II Web Edition software supports the entire Cyclone III family. Quartus II Web Edition software offers productivity and performance features, including TimeQuest timing analysis, PowerPlay power optimization, and SOPC Builder, providing the fastest path to design completion for Cyclone III FPGAs.

**90-nm low-cost FPGAs**

As the second-generation offering in the Cyclone series, the Cyclone II FPGA family offers more density and features at dramatically lower costs, compared with its predecessor. Based on a 1.2V, 90-nm, low-k dielectric process, the devices include dedicated DSP circuitry for very low-cost DSP solutions. If you're implementing Nios II embedded processors on Cyclone II FPGAs, you can create cost-effective processing solutions for price-sensitive and compute-intensive applications.

Cyclone II family features summary

Cost-optimized architecture	Offers from 4,608 to 68,416 LEs—3.5X the density of first-generation Cyclone FPGAs—and the lowest cost per LE.
Embedded memory	Up to 1.1 Mbits of RAM via 4,608-bit memory blocks capable of 250-MHz performance. Support for multiple configurations, including true dual-port and single-port RAM, ROM, and FIFO buffers.
Embedded multipliers	150 18-bit x 18-bit embedded multipliers running at 250 MHz can implement common DSP functions such as finite impulse response (FIR) filters and fast Fourier transforms (FFTs). Each 18-bit x 18-bit multiplier can be configured as two independent 9-bit x 9-bit multipliers.
Nios II embedded processor support	Nios II embedded processors offer an ideal replacement for low-cost discrete microprocessors to reduce cost and increase flexibility. Nios II soft processor support for Cyclone II FPGAs offers over 100-DMIPS performance.
Differential and single-ended I/O standards support	Support for LVTTTL, LVCMOS, PCI, PCI-X, PCI Express ¹ , SSTL, and high-speed transceiver logic (HSTL) single-ended I/O standards. Differential signaling support for LVDS (805-Mbps receiving and 640-Mbps transmitting), mini-LVDS, RSDS, LVPECL, SSTL, and HSTL system interfaces.
External memory interfaces	Dedicated interfaces supporting external memory devices at 167 MHz for integration with external SDR, DDR, DDR2 SDRAM, and QDR II SRAM devices. Altera offers DDR, DDR2, and QDR II memory controller MegaCore functions free with Quartus II software subscriptions.
Clock management circuitry	16 low-skew, global clock networks span the entire device, fed by 16 dedicated input clock pins. Four PLLs provide complete system clock management on and off chip. Each PLL has three output taps, and features programmable bandwidth, programmable duty cycle, spread-spectrum clocking, lock detection, and frequency synthesis with phase-shifting capabilities.
On-chip termination (OCT)	Single-ended OCT support for driver impedance matching and series termination eliminates the need for external resistors, improves signal integrity, and simplifies board design.
Hot socketing and power sequencing	Robust on-chip hot-socketing and power-sequencing support ensures proper device operation independent of the power-up sequence.
Automatic SEU detection circuitry	Features automatic SEU detection circuitry utilizing 32-bit CRC.

¹ Requires external PHY device



130-nm low-cost FPGAs

Cyclone devices provide application-focused features such as embedded memory, external memory interfaces, and clock management circuitry at price points optimal for high-volume applications.

Cyclone family features summary

Cost-optimized architecture	Offers from 2,910 to 20,060 LEs; built for low cost.
Embedded memory	288 Kbits of RAM through 4,608-bit memory blocks can be configured to support a wide range of operation modes, including RAM, ROM, FIFO buffers, and single-port and dual-port modes.
Nios II embedded processor support	Nios II embedded processors offer an ideal replacement for low-cost discrete microprocessors to reduce cost and increase flexibility.
Single-ended I/O support	Supports a variety of single-ended I/O interface standards, such as the 3.3V, 2.5V, 1.8V, LVTTTL, LVCMOS, SSTL, and PCI standards needed for today's systems.
External memory interfaces	Dedicated external memory interfaces allow you to integrate external SDR and DDR SDRAM devices into complex system designs without degrading data access performance.
Differential I/O support	Support for 129 LVDS and RSDS channels with 640-Mbps LVDS data rates and 311-Mbps RSDS data rates.
Clock management circuitry	Features up to two programmable PLLs and eight global clock lines that provide robust clock management and frequency synthesis capabilities enabling on- and off-chip system clock management. PLLs offer advanced features such as frequency synthesis, programmable phase shift, programmable delays, and external clock output.
Hot socketing and power sequencing	Robust on-chip hot-socketing and power-sequencing support that ensures that Cyclone devices operate no matter how the system is powered up.
Automatic SEU detection circuitry	Utilizes 32-bit CRC to minimize radiation problems.

Cyclone FPGA series features

		Cyclone III family (1.2V)							
		EP3C5	EP3C10	EP3C16	EP3C25	EP3C40	EP3C55	EP3C80	EP3C120
Density and speed	LEs	5,136	10,320	15,408	24,624	39,600	55,856	81,264	119,088
	Total RAM (Kbits) ¹	414	414	504	594	1,134	2,340	2,745	3,888
	M9K block (8 Kbits + 512 parity bits) ¹	46	46	56	66	126	260	305	432
	Speed grades (fastest to slowest) ²	6, 7, 8	6, 7, 8	6, 7, 8	6, 7, 8	6, 7, 8	6, 7, 8	6, 7, 8	7, 8
Architectural features	Embedded processor	Nios II processor							
	18-bit x 18-bit/9-bit x 9-bit embedded multipliers	23/46	23/46	56/112	66/132	126/252	156/312	244/488	288/576
	True dual-port RAM	✓	✓	✓	✓	✓	✓	✓	✓
	Global clock networks	10	10	20	20	20	20	20	20
	PLLs/unique outputs	2/10	2/10	4/20	4/20	4/20	4/20	4/20	4/20
	Configuration file size (Mbits)	2.8	2.8	3.9	5.5	9.1	14.2	19	27.2
I/O features	I/O voltage levels supported	1.5, 1.8, 2.5, 3.0, 3.3							
	I/O standards supported	LVDS, RSDS, Mini-LVDS, LVPECL, Differential SSTL-18 (I and II), Differential SSTL-2 (I and II), 1.5V Differential HSTL (I and II), 1.8V Differential HSTL (I and II), SSTL-18 (I and II), SSTL-2 (I and II), 1.5V HSTL (I and II), PCI, PCI-X, PCI Express ³ , LVTTL, LVCMOS, PPDS							
	LVDS maximum data rate (Mbps) (receive/transmit)	875/840							
	LVDS channels	66	66	136	79	223	159	177	229
	RSDS maximum data rate (Mbps) (transmit)	360							
	Mini-LVDS maximum data rate (Mbps) (transmit)	400							
	Series OCT	✓	✓	✓	✓	✓	✓	✓	✓
	Programmable drive strength	✓	✓	✓	✓	✓	✓	✓	✓
External memory interfaces	Memory device supported	QDR II, DDR2, DDR, SDR							
	MegaCore controller with clear text datapath	✓							
	System timing and analysis	✓							
	Board layout guideline	✓							

¹ Kbits = 1,024 bits² Not all packages are supported in all speed grades³ Requires external PHY device

Cyclone FPGA series features (continued)

		Cyclone II family (1.2V)							Cyclone family (1.5V)				
		EP2C5/A	EP2C8/A	EP2C15A	EP2C20/A	EP2C35	EP2C50	EP2C70	EP1C3	EP1C4	EP1C6	EP1C12	EP1C20
Density and speed	LEs	4,608	8,256	14,448	18,752	33,216	50,528	68,416	2,910	4,000	5,980	12,060	20,060
	Total RAM (Kbits) ¹	117	162	234	234	472.5	580.5	1,125	58.5	76.5	90	234	288
	M4K RAM blocks (4 Kbits ¹ + 512 parity bits)	26	36	52	52	105	129	250	13	17	20	52	64
	Speed grades (fastest to slowest)	-6, -7, -8							-6, -7, -8				
Architectural features	Embedded processor available	Nios II processor							Nios II processor				
	18-bit x 18-bit/9-bit x 9-bit embedded multipliers	13/26	18/36	26/52	26/52	35/70	86/172	150/300	–	–	–	–	–
	True dual-port RAM	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	Global and regional clock networks	8	8	16	16	16	16	16	8	8	8	8	8
	PLLs/unique outputs	2/6	2/6	4/12	4/12	4/12	4/12	4/12	1/3	2/6	2/6	2/6	2/6
	Configuration file size (Mbits)	1.26	1.98	3.89	3.89	6.85	9.96	14.31	0.63	0.93	1.17	2.32	3.56
I/O features	I/O voltage levels supported (V)	1.5, 1.8, 2.5, 3.3							1.5, 1.8, 2.5, 3.3				
	I/O standards supported	LVDS, RSDS, Mini-LVDS, LVPECL, Differential SSTL-18 (I and II), Differential SSTL-2 (I and II), 1.5V Differential HSTL (I and II), 1.8V Differential HSTL (I and II), SSTL-18 (I and II), SSTL-2 (I and II), 1.5V HSTL (I and II), PCI, PCI-X, PCI Express ² , LVTTTL, LVCMOS							LVDS, RSDS, Differential SSTL-2, SSTL-2 (I and II), SSTL-3 (I and II), PCI, LVTTTL, LVCMOS				
	LVDS maximum data rate (Mbps) (receive/transmit)	805/640	805/640	805/640	805/640	805/640	805/640	805/640	640/640	640/640	640/640	640/640	640/640
	LVDS channels	60	79	136	136	209	197	265	34	129	72	103	129
	RSDS maximum data rate (Mbps) (transmit)	311	311	311	311	311	311	311	311	311	311	311	311
	Mini-LVDS maximum data rate (Mbps) (transmit)	311	311	311	311	311	311	311	–	–	–	–	–
	Series OCT	✓	✓	✓	✓	✓	✓	✓	–	–	–	–	–
	Programmable drive strength	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
External memory interfaces	Memory devices supported	QDR II, DDR2, DDR, SDR							DDR, SDR				
	MegaCore controller with clear text datapath	✓							✓				
	System timing analysis	✓							✓				
	Board layout guidelines	✓							✓				

¹ Kbits = 1,024 bits

² Requires external PHY device

Cyclone FPGA series comparative table

Device	LEs	Embedded RAM blocks ¹	RAM Kbits	18-bit x 18-bit multipliers	PLLs	Clock outputs per PLL	I/O banks	Max LVDS channels ²
EP1C3	2,910	13	58	-	1	3	4	34
EP1C4	4,000	17	76	-	2	3	4	129
EP2C5	4,608	26	119	13	2	3	4	61
EP3C5	5,136	46	414	23	2	5	8	62
EP1C6	5,980	20	90	-	2	3	4	72
EP2C8/A	8,256	36	165	18	2	3	4	75
EP3C10	10,320	46	414	23	2	5	8	62
EP1C12	12,060	52	234	-	2	3	4	103
EP2C15A	14,448	52	239	26	4	3	8	128
EP3C16	15,408	56	504	56	2	5	8	128
EP1C20	20,000	64	288	-	4	3	4	129
EP2C20/A	18,752	52	239	35	4	3	8	128
EP3C25	24,624	66	594	66	4	5	8	71
EP2C35	33,216	105	483	35	4	3	8	201
EP3C40	39,600	126	1,134	126	4	5	8	215
EP2C50	50,528	129	594	86	4	3	8	189
EP3C55	55,856	260	2,340	156	4	5	8	151
EP2C70	68,416	250	1,152	150	4	3	8	257
EP3C80	81,264	305	2,745	244	4	5	8	169
EP3C120	119,088	432	3,888	288	4	5	8	221

¹ Cyclone and Cyclone II FPGAs have 4-Kbit RAM blocks; Cyclone III FPGAs have 9-Kbit RAM blocks.

² Bidirectional I/O pins can be used as inputs or outputs; does not include dedicated clock input pins.

Cyclone FPGA series package matrix and maximum user I/Os

Device	Package/package size (mm)												
	T100 16 x 16	E/T1443 22 x 22	Q208 30 x 30	Q240 35 x 35	F256 17 x 17	U256 14 x 14	F324 19 x 19	F400 21 x 21	F484 23 x 23	U484 19 x 19	F672 27 x 27	F780 29 x 29	F896 31 x 31
EP1C3	65	104											
EP1C4							249	301					
EP2C5		89	142		158								
EP3C5		94			182	182							
EP1C6		98		185	185								
EP2C8/A		85	138		182								
EP3C10		94			182	182							
EP1C12				173	185		249						
EP2C15A					152				315				
EP3C16		84		160	168	168			346	346			
EP1C20							233	301					
EP2C20/A				142	152				315				
EP3C25		82		148	156	156	215						
EP2C35									322	322	475		
EP3C40				128			195		331	331		535	
EP2C50									294	294	450		
EP3C55									327	327		377	
EP2C70											422		622
EP3C80									295	295		429	
EP3C120									283			531	

FOR MORE INFORMATION

Cyclone III FPGAs www.altera.com/cyclone3
 Cyclone II FPGAs www.altera.com/cyclone2
 Cyclone FPGAs www.altera.com/cyclone
 Training www.altera.com/training



Transceiver FPGAs optimized for mainstream applications up to 3.125 Gbps

With our Arria® GX FPGAs, you can quickly and easily connect custom logic to devices with transceiver speeds up to 3.125 Gbps. Arria GX devices use proven transceiver technology that supports a variety of serial protocols, allowing you to bridge legacy parallel interfaces to serial interfaces or to bridge serial interfaces to serial interfaces. This family offers logic densities ranging from 20K to 90K LEs and supports a number of ready-made IP cores to get your designs up and running.

Arria GX FPGAs are backed by years of Altera® transceiver expertise: a physical coding sublayer (PCS) and physical medium attachment (PMA) interface are protocol-optimized, low-power versions of those in Stratix® II GX FPGAs. Arria GX devices give you an easy way to interface custom logic to mainstream protocols at a very low total cost of ownership. You'll also benefit from Altera's extensive set of reliable tools and comprehensive support infrastructure.

Arria GX features summary

Complete solution	Fully tested IP is available for PCI Express (x1 and x4), Gigabit Ethernet, serial digital interface (SDI), and Serial RapidIO® up to 3.125 Gbps. A simple and inexpensive development kit with reference designs is also available to demonstrate the simplicity of designing transceiver interfaces with Arria GX devices.
Best-in-class signal integrity	The Arria GX transceiver is based on the proven Stratix II GX transceiver and built with flip chip packages. For those protocols supported by Arria GX FPGAs, you'll get signal integrity that's close to what can be achieved by a Stratix II GX FPGA.
Unparalleled software tools	The award-winning Quartus® II Web Edition, as well as the subscription version, supports all members of the Arria GX family. This makes designing with this family easy and inexpensive.
Extensive support network	Altera's team of several hundred field applications engineers (FAEs) worldwide are available to help drive your success. This team is also backed by a global team of applications engineers in several Regional Service Centers around the world.

Arria GX transceiver protocol support

Protocol	Data rate
PCI Express 1.1	2.5 Gbps
XAUI	3.125 Gbps
Serial RapidIO	1.25, 2.5, 3.125 Gbps
Gigabit Ethernet	1.25 Gbps
Fibre Channel FC-2, FC-1	1.0625, 2.125 Gbps
SDI 3G	3 Gbps
SDI, standard definition (SD), high definition (HD)	0.271, 1.5 Gbps
SerialLite II	Up to 3.125 Gbps
Basic mode	Up to 3.125 Gbps
Common Public Radio Interface (CPRI)	2.45 and 3.072 Gbps
Open Base Station Architecture Initiative (OBSAI)	0.768, 1.536, and 3.072 Gbps

Arria GX FPGA package and I/O matrix

36 Number indicates available user I/O pins.

Vertical migration (Same V_{CC}, GND, ISP, and input pins).

All Arria GX devices are offered in commercial and industrial temperatures and RoHS-compliant packages.

Arria GX FPGAs (1.2V) Low cost, risk free

		EP1AGX20	EP1AGX35	EP1AGX50	EP1AGX60	EP1AGX90
FineLine BGA (F)	484 pin (flip chip)	255	255	254	254	
	780 pin	366	366	395	395	
	1,152 pin			563	583	607

Arria GX FPGA family features

		Arria GX FPGAs (1.2V) Low cost, risk free				
		EP1AGX20	EP1AGX35	EP1AGX50	EP1AGX60	EP1AGX90
Density and speed	Equivalent LEs	21,580	33,520	50,160	60,100	90,220
	Adaptive logic modules (ALMs)	8,632	13,408	20,064	24,040	36,088
	Adaptive look-up tables (ALUTs)	17,264	26,816	40,128	48,080	72,176
	Total RAM (Kbits) ¹	1,229	1,348	2,475	2,529	4,478
	M512 RAM blocks (512 bits + 64 parity bits)	166	197	313	326	478
	M4K RAM blocks (4 Kbits ¹ + 512 parity bits)	118	140	242	252	400
	M-RAM blocks (512 Kbits ¹ + 65,536 parity bits)	1	1	2	2	4
	18-bit x 18-bit/9-bit x 9-bit embedded multipliers	40/80	56/112	104/208	128/256	176/352
	Speed grades (fastest to slowest)	-6	-6	-6	-6	-6
Architectural features	Embedded processor available	Nios® II processor				
	DSP blocks	10	14	26	32	44
	I/O registers per I/O element	6	6	6	6	6
	True dual-port RAM	✓	✓	✓	✓	✓
	Global and regional clock networks	48	48	48	48	48
	PLLs/unique outputs	4	4	4 or 8	4 or 8	8
	Design security	—	—	—	—	—
	HardCopy® II device support	—	—	—	—	—
I/O features	I/O voltage levels supported (V)	1.5, 1.8, 2.5, 3.3				
	I/O standards supported	LVDS, LVPECL, Differential SSTL-18, Differential SSTL-2, Differential HSTL, SSTL-18 (I and II), SSTL-2 (I and II), 1.5V HSTL (I and II), 1.8V HSTL (I and II), PCI, PCI-X 1.0, LVTTTL, LVCMOS				
	True-LVDS maximum data rate (Mbps)	125–840				
	True-LVDS channels (receive/transmit)	31/29	31/29	42/42	42/42	47/45
	Embedded DPA circuitry	✓	✓	✓	✓	✓
	Series and differential OCT	✓	✓	✓	✓	✓
	Programmable drive strength	—	—	—	—	—
	Transceiver (SERDES) data rate range	600 Mbps to 3.125 Gbps				
	Transceiver (SERDES) channels	4	4 or 8	4 or 8	4, 8 or 12	12
External memory interfaces	Memory devices supported	DDR2, DDR, SDR				
Configuration file sizes	Configuration file size (Mbits)	10	10	17	17	28

¹ Kbits = 1,024 bits

Stratix high-end FPGA series

The Stratix® FPGA series enables you to deliver high-performance, state-of-the-art products to market faster with lower risk and higher productivity. By combining high density, high performance, and a rich feature set, the Stratix FPGA series allows you to integrate more functions and maximize performance. Using the Quartus® II design software suite, along with a broad portfolio of IP, you can gain the highest level of productivity for large and complex team-based designs. Also contributing to the total solution are HardCopy® ASICs, which provide a seamless migration path to low-cost volume production. Backed by Altera's track record as a reliable and high-quality supplier, the Stratix FPGA series provides the tools you need to design with confidence.

Key features

- Industry's biggest and fastest FPGAs
- Lowest power high-end FPGAs
- Flexible and high-performance I/O pins
- Highest DSP and memory capabilities
- Path to low-cost HardCopy ASICs
- Volatile and non-volatile design security
- Quartus II design software for highest performance and productivity



40-nm high-performance, high-end FPGAs with 8.5-Gbps transceiver option

The Stratix IV FPGA family gives you the highest density, highest system performance, and the lowest power available in an FPGA, enabling innovation without compromise. Two variants meet the diverse application needs of large designs: Stratix IV GX FPGAs and Stratix IV E FPGAs. Leveraging the benefits of the 40-nm process and proven transceiver and memory interface technology, the Stratix IV GX device delivers unprecedented system bandwidth with superior signal

integrity. Both Stratix IV GX and Stratix IV E FPGA variants include the industry's most efficient and highest-performance logic, embedded memory, and DSP capabilities. Additionally, both device variants deliver a seamless low-risk path to production volumes with our HardCopy IV ASICs. Stratix IV FPGAs are your best total solution for high-end system on a chip (SOC) design.

Stratix IV family features summary (GX only)

Multi-gigabit transceivers (PMA and PCS features)	Up to 48 high-speed transceiver channels (full duplex) with PLL-based clock data recovery (CDR) operating from 600 Mbps to 8.5 Gbps (6.375 Gbps industrial) and superior signal integrity. Up to 32 channels at 8.5 Gbps, plus up to 16 additional channels at 3.2 Gbps.
Enhanced transceiver block	Two central clock management units (CMUs) per transceiver block, used for reference clock inputs and transmit PLLs. CMUs can be configured as transmit PLLs or as additional channels 5 and 6 in the transceiver block. Embedded PCS pattern detector, word aligner with programmable pattern byte reordering, 8b/10b encoder/decoder, protocol-specific hardware for PCI Express (PIPE interface), XAUI, Gigabit Ethernet, and SONET.
Protocol support	Support for PCI Express Gen1 and Gen2, XAUI/HiGig-XAUI, Gigabit Ethernet, Serial RapidIO®, Interlaken, CEI-6G, SONET OC-3/12/48, HyperTransport™ 3.0, SDI, ASI, GPON, SATA/SAS, CPRI, OBSAI, and 3G and 6G basic modes with rule-based protocol configuration.
Plug & Play Signal Integrity	Support for adaptive equalization and hot socketing.
PCI Express hard IP	Complete protocol solution with embedded PCI Express hard IP blocks, Gen1 and Gen2, that implement the PHY-MAC layer, data-link layer, and transaction layer functionality.
Dynamic transceiver reconfiguration	Provides the capability to support multiple protocols and data rates on the same transceiver channel without reprogramming the FPGA.
Low-power transceiver	Typical PMA power consumption of 100 mW at 3.125 Gbps and 135 mW at 6.375 Gbps per channel.

Stratix IV family features summary (E and GX)

ALMs	With up to 530K equivalent LEs on the GX variant and up to 680K equivalent LEs on the E variant, Stratix IV ALMs provide industry-leading performance and logic density with a patented fracturable eight-input LUT, two dedicated adders, and two registers. ALMs deliver 80 percent greater logic density than competitive architectures along with more registers, resulting in increased performance through reduced total logic levels and associated routing.
Performance	Fastest FPGA in the industry, through a high-speed fabric plus the most efficient routing architecture, as well as the industry's most efficient embedded memory and DSP architecture.
Programmable Power Technology	Every programmable logic array block (LAB), DSP block, and memory block delivers high-speed or low power, depending on your design requirements. Using the Programmable Power Technology in Stratix IV FPGAs, all functional blocks in the array, except those designated as timing critical, are set to low-power mode. With only the timing-critical blocks set to high-speed mode, power dissipation in Stratix IV devices is reduced by up to 50 percent.
External memory interface circuitry	Stratix IV FPGAs provide the industry's highest performance external DRAM and SRAM memory interfaces, including DDR3, DDR2, QDR II, QDR, RLDRAM II, and RLDRAM. The dynamic self-calibrating PHY continuously compensates for PVT variations during operation. You can achieve higher performance, up to 1,067 Mbps (533 MHz), better signal integrity, increased robustness, expanded design margin, and lower power in your designs. Complete solutions including easy-to-use IP, hardware development platforms, and comprehensive timing analysis are available to minimize risk and speed time to market.
HardCopy ASIC series support	The HardCopy ASIC migration path is now available for both Stratix IV GX and E variants, with a single toolset across Stratix IV and HardCopy IV devices for enhanced productivity. All of the Stratix IV devices and the HardCopy IV ASICs are supported in Quartus II software.
MultiTrack interconnect	Connections between ALMs, TriMatrix memory blocks, DSP blocks, and device I/O pins are provided by the MultiTrack interconnect structure. Compared to competitors, the MultiTrack interconnect offers nearly 3X the number of one-hop interconnects between LABs, maximizing performance for challenging, time-critical logic. The interconnect has the industry's largest single-hop span of five LABs, minimizing routing congestion and further improving performance and routability. The Quartus II compiler automatically places critical design paths on faster interconnects to maximize performance.
TriMatrix memory	TriMatrix memory offers three memory block sizes to suit complex design needs: 640-bit memory logic array blocks (MLABs), 9-Kbit M9K blocks, and 144-Kbit M144K blocks. With up to 23 Mbits of memory performing at over 600 MHz and features such as dual-port RAM and integrated error correction code (ECC), TriMatrix memory is more flexible and efficient, and provides higher memory bandwidth than any other memory architecture. All three configurations support parity checking.
DSP blocks	Stratix IV DSP blocks provide more than twice as many multiplier resources as competing architectures, delivering the industry's fastest performance with up to 1,360 18-bit by 18-bit multipliers. DSP blocks include all associated pipelining, adders, and accumulators. They are configurable to support 9-bit x 9-bit, 12-bit x 12-bit, 18-bit x 18-bit, or 36-bit x 36-bit multipliers at up to 550 MHz and 748 GMACS, making Stratix IV FPGAs ideal for custom-built DSP replacement, as well as coprocessing for hardware acceleration.
High-speed single-ended I/O support	Up to 24 modular I/O banks with programmable slew rate, programmable drive strength, programmable output delay, OCT, and dynamic trace compensation. Maximizes I/O flexibility and performance for over 40 industry standards to optimally match your system requirements.
High-speed differential I/O support	Optimized LVDS I/O pins at between 150-MHz and 1.6-GHz performance provide high performance and excellent signal integrity with half the capacitance of competitive devices. Stratix IV LVDS supports hard dynamic phase alignment (DPA) and serializer/deserializer (SERDES) blocks with clock forwarding for CDR in high-speed chip-to-chip applications. MegaWizard® Plug-Ins make differential I/O configuration straightforward and easy.
DPA	As a feature of the LVDS high-speed differential I/O support, DPA minimizes bit errors and simplifies PCB layout and timing management for high-speed data transfer. DPA eliminates channel-to-channel and channel-to-clock skew in high-speed data transmission systems.
OCT	Serial, parallel, dynamic, and differential OCT simplifies design and reduces component count. Dynamic OCT enables support for next-generation memory interface designs that must dynamically switch between serial and parallel termination. Digital calibration circuitry provides industry-leading OCT tolerance within 10 percent.
Clock management features	Each product family has up to 12 PLLs with up to 10 unique customizable outputs per PLL ranging from 5 to 720 MHz. There are also up to 16 global, 88 quadrant, and 208 periphery clocks. Advanced clock management features include PLL reconfiguration, frequency synthesis, programmable phase shift, programmable delay shift, external feedback, and programmable bandwidth for all on- and off-chip clocking requirements.
Design security	Volatile and non-volatile design security protects FPGA design and IP against copying, reverse engineering, and tampering.
Remote system upgrades	This capability enables reliable and safe deployment of in-system enhancements and bug fixes.
Hot socketing and power sequencing	Robust on-chip hot-socketing and power-sequencing support ensures proper device operation independent of the power-up sequence.
Automatic SEU detection circuitry	Automatic SEU detection circuitry utilizes 32-bit CRC with criticality processor to determine the importance of error.
Nios® II embedded processor support	Nios II embedded processors reduce cost, increase flexibility, and offer an ideal replacement for low-cost discrete microprocessors. Nios II soft processor support for Stratix IV FPGAs offers over 300-DMIPS performance.

Stratix IV transceiver protocol support (GX only)

Protocol	Data rate	Protocol	Data rate
PCI Express Gen1, Gen2	2.5, 5.0	DDR-XAUI	6.25
PCI Express Cable	2.5	SF15.1	2.488 - 3.125
SAS	1.5, 3.0, 6.0	SONET OC-3/OC-12/OC-48	0.155, 0.622, 2.488
SATA	1.5, 3.0	Ethernet-XAUI	4 x 3.125
SDI SD/HD	0.27/1.485	Gigabit Ethernet	1.25
3G SDI	2.97	CEI-6G/SR/LR	6.375/4.976 - 6.375/4.976 - 6.375
ASI	0.27	HiGig+	3.75
CPRI	0.6144, 1.2288, 2.4576, 3.072	HiGig2	4.0625
OBSAI	0.768, 1.536, 3.072	Interlaken	3.125 - 6.375
Serial RapidIO	1.25, 2.5, 3.125	Fibre Channel	1.0625, 2.125, 4.25, 8.5
GPON	1.244 uplink, 2.488 downlink	HyperTransport 3.0	0.4, 2.4, 2.8, 3.2
SPAUI	3.2, 6.4	SerialLite II	0.6 - 6.375

**65-nm high-performance, high-end FPGAs**

Altera's Stratix III FPGAs were designed to address your design and business challenges, anticipate the unforeseen, and help you win in your market. The Stratix III family supports your applications' increased levels of integration and complexity with higher densities and performance, while decreasing power consumption. Its flexible and efficient logic architecture, enhanced memory blocks, and high-capacity DSP blocks meet your system's most demanding requirements.

Specifically designed for ease of use and rapid system integration, the Stratix III FPGA family offers two variants optimized to meet different application needs. Stratix III L FPGAs deliver balanced logic, memory, and DSP resources for general-purpose applications; Stratix III E FPGAs provide enhanced memory and DSP resources for memory- and DSP-intensive applications. Both variants include fast, flexible, and robust I/O connectivity, providing DDR3 DIMM memory interface support up to 1,067 Mbps.

Stratix III family features summary

HardCopy ASIC series support	This pin-compatible HardCopy ASIC reduces cost and power for volume applications.
Programmable Power Technology	The core logic, routing, memory, and DSP blocks in Stratix III FPGAs have fine-grained control over whether a specific tile is in high-speed mode or low-power mode. Quartus II design software automatically applies the high-speed mode to all parts of the design that require the highest performance, utilizing the low-power mode for all other parts of the FPGA. Leakage power is reduced by about 70 percent for the portions of the design implemented with low-power mode.
Selectable core voltage	The core voltage of Stratix III FPGAs can be set depending on performance requirements—to the 1.1V setting for the highest performance and 0.9V for the lowest power. The 0.9V setting reduces dynamic power by 55 percent over previous-generation 1.2V devices. The core voltage setting is independent of Programmable Power Technology.
ALMs	ALMs provide industry-leading performance and logic density with a patented fracturable eight-input LUT, two dedicated adders, and two registers. ALMs deliver 80 percent greater logic density than competitive architectures along with more registers, resulting in increased performance through reduced total logic levels and associated routing.
MultiTrack interconnect	Connections between ALMs, TriMatrix memory blocks, DSP blocks, and device I/O pins are provided by the MultiTrack interconnect structure. Compared to competitors, the MultiTrack interconnect offers nearly 3X the number of one-hop interconnects between LABs, maximizing performance for challenging, time-critical logic. The interconnect has the industry's largest single-hop span of five LABs, minimizing routing congestion and further improving performance and routability. The Quartus II compiler automatically places critical design paths on faster interconnects to maximize performance.
TriMatrix memory	TriMatrix memory offers three memory block sizes to suit complex design needs: 640-bit MLABs, 9-Kbit M9K blocks, and 144-Kbit M144K blocks. With up to 21 Mbits of memory performing at over 600 MHz and features such as dual-port RAM and integrated ECC, TriMatrix memory is more flexible and efficient, and provides higher memory bandwidth than any other memory architecture. All three configurations support parity checking.
DSP blocks	DSP blocks provide twice as many multiplier resources as competing architectures, with up to 896 18-bit x 18-bit multipliers on the EP3SE110 device. DSP blocks include all associated pipelining, adders, and accumulators. They are configurable to support 9-bit x 9-bit, 12-bit x 12-bit, 18-bit x 18-bit, or 36-bit x 36-bit multipliers at up to 550 MHz.
External memory interface circuitry	Stratix III FPGAs provide the industry's highest performance external DRAM and SRAM memory interfaces, including DDR3, up to 1,067-Mbps DIMM memory interface support, PVT-compensated dedicated hard I/O structures with OCT, trace compensation, read/write leveling, and half- and full-rate registered outputs. Complete solutions, including easy-to-use IP, hardware development platforms, and comprehensive timing analysis, are available to minimize risk and speed time to market.
High-speed single-ended I/O support	Up to 24 modular I/O banks with dedicated DQ/DQS circuitry, programmable slew rate, programmable drive strength, programmable output delay, OCT, and dynamic trace compensation maximize I/O flexibility and performance for over 40 industry standards to optimally match your system requirements.
High-speed differential I/O support	Optimized LVDS I/O pins provide high performance and excellent signal integrity with half the capacitance of competitive devices. Stratix III LVDS supports hard DPA and SERDES blocks with clock forwarding for clock data recovery in high-speed chip-to-chip applications. Wizards make differential I/O configuration straightforward and easy.
DPA	DPA minimizes bit errors and simplifies PCB layout and timing management for high-speed data transfer. DPA eliminates channel-to-channel and channel-to-clock skew in high-speed data transmission systems.
OCT	Serial, parallel, dynamic, and differential OCT simplifies design and reduces component count. Dynamic OCT enables support for next-generation memory interface designs that must dynamically switch between serial and parallel termination. Digital calibration circuitry provides industry-leading OCT tolerance within 10 percent.
Clock management features	Each product family has up to 12 PLLs with up to 10 unique customizable outputs per PLL ranging from 5 to 720 MHz. There are also up to 16 global, 88 quadrant, and 208 periphery clocks. Advanced clock management features include PLL reconfiguration, frequency synthesis, programmable phase shift, programmable delay shift, external feedback, and programmable bandwidth for all on- and off-chip clocking requirements.
Design security	To prevent IP theft and product tampering, the devices feature configuration bitstream encryption using the Advanced Encryption Standard (AES) and 256-bit key. Stratix III FPGAs are the industry's only design security solution with both non-volatile and volatile key options. The non-volatile key is more practical, easy to use, and low cost, while the volatile key is reprogrammable.
Remote system upgrades	This capability enables reliable and safe deployment of in-system enhancements and bug fixes.
Hot socketing and power sequencing	Robust on-chip hot-socketing and power-sequencing support ensures proper device operation independent of the power-up sequence.
Automatic SEU detection circuitry	The devices feature automatic SEU detection circuitry utilizing 32-bit CRC with criticality processor to determine the importance of error.
Nios II embedded processor support	Nios II embedded processors reduce cost, increase flexibility, and offer an ideal replacement for low-cost discrete microprocessors. Nios II soft processor support for Stratix III FPGAs offers over 300-DMIPS performance.



90-nm high-performance, high-density FPGAs with 6.375-Gbps transceiver option

Stratix II devices are the industry's leading 90-nm high-density, high-performance FPGAs. Built on an ALM logic structure to maximize performance and minimize power, Stratix II devices on average deliver 50 percent faster performance than prior-generation FPGAs. The Stratix II device family has variants both with and without transceivers. The Stratix II GX devices are the family's transceiver variant with 20 low-power, high-speed channels and data rates from 600 Mbps to

6.375 Gbps. Dynamically configurable transmit pre-emphasis and receiver equalization optimizes signal integrity under adverse channel conditions. The transceivers are capable of driving FR-4 backplanes at 6.375 Gbps and have proven to be interoperable with backplanes and transceivers from multiple vendors. In addition, the PCS hard IP, which is part of the transceiver block, saves valuable logic resources and simplifies protocol support.

Stratix II GX features summary

Excellent signal integrity, 600 Mbps to 6.375 Gbps	The transmitter has low jitter generation and up to 500 percent pre-emphasis. The receiver has excellent jitter tolerance and up to 17-dB of equalization.
Low power	The transceiver (including both PMA and PCS blocks) dissipates 200 mW per channel at 6.375 Gbps and only 140 mW per channel at 3.125 Gbps.
PCS support (hard IP)	The transceiver supports the following PCS blocks: PCI Express, PIPE-compliant PCS, CEI-6G, 8b/10b encoder/decoder, XAUI state machine and channel bonding, Gigabit Ethernet state machine, SONET, and 8/10/16/20/32/40-bit interface (to logic).
System-level diagnostics	Serial loopback, reverse serial loopback, psuedo-random binary sequence (PRBS) generator checker, and register-based interface facilitate dynamic reconfiguration of pre-emphasis, equalization, and differential output voltage.

Stratix II family features summary (with and without GX transceivers)

HardCopy II ASIC support	This pin-compatible HardCopy ASIC reduces cost and power for volume applications. (Not available for Stratix II GX transceiver variant.)
ALMs	ALMs provide industry-leading performance and logic density with a patented fracturable eight-input LUT, two dedicated adders, and two registers. ALMs deliver 80 percent greater logic density than competitive architectures along with more registers, resulting in increased performance through reduced total logic levels and associated routing.
MultiTrack interconnect	Connections between ALMs, TriMatrix memory blocks, DSP blocks, and device I/O pins are provided by the MultiTrack interconnect structure. Compared to competitors, the MultiTrack interconnect offers nearly 3X the number of one-hop interconnects between LABs, maximizing performance for challenging, time-critical logic. The interconnect has the industry's largest single-hop span of five LABs, minimizing routing congestion and further improving performance and routability. The Quartus II compiler automatically places critical design paths on faster interconnects to maximize performance.
TriMatrix memory	TriMatrix memory in Stratix II FPGAs provides up to 9 Mbits of memory in three block sizes: 512 bits in the M512, 4 Kbits in the M4K, and 500 Kbits in the M-RAM blocks. TriMatrix memory includes parity checking and is capable of up to 550-MHz performance.
DSP blocks	Each DSP block includes multipliers with associated pipelining, adders, and accumulators. Stratix II FPGAs include up to 96 DSP blocks offering 384 18-bit x 18-bit multipliers that operate at up to 450 MHz. Each multiplier can be configured as 9-bit x 9-bit, 18-bit x 18-bit, or 36-bit x 36-bit multipliers.
External memory interface circuitry	Fully characterized and hardware-validated circuitry provides flexible, high-performance interfaces to the latest external DRAM and SRAM memory technologies. Complete solutions, including easy-to-use IP, hardware development platforms, and comprehensive timing analysis, are available to minimize risk and speed time to market.
High-speed differential I/O support	Optimized LVDS I/Os with DPA and SERDES capability provides high performance and excellent signal integrity. Wizards make differential I/O configuration straightforward and easy.
DPA	DPA maximizes signal integrity and simplifies PCB layout and timing management for high-speed data transfer. DPA eliminates channel-to-channel and channel-to-clock skew in high-speed data transmission systems.

Stratix II family features summary (with and without GX transceivers) (continued)

OCT	Serial and differential OCT simplifies design and reduces component count. Digital calibration circuitry provides industry-leading OCT tolerance within five percent.
Clock management features	Up to 12 PLLs and up to 48 system clocks. Advanced clock management features include PLL reconfiguration, spread-spectrum clocking, frequency synthesis, programmable phase shift, programmable delay shift, external feedback, and programmable bandwidth for all on- and off-chip clocking requirements.
Design security	Configuration bitstream encryption using 128-bit AES prevents IP theft and product tampering. The non-volatile key storage makes the solution reliable, easy to use, and low cost.
Remote system upgrades	Enables reliable and safe deployment of in-system enhancements and bug fixes.
Hot socketing and power sequencing	Robust on-chip hot-socketing and power-sequencing support ensures proper device operation independent of the power-up sequence.
Automatic SEU detection circuitry	Features automatic SEU detection circuitry utilizing 32-bit CRC.
Nios II embedded processor support	Nios II embedded processors reduce cost, increase flexibility, and offer an ideal replacement for low-cost discrete microprocessors. Nios II soft processor support for Stratix II FPGAs offers over 200-DMIPS performance.

Stratix II GX protocol support

Protocol	Data rate
3G SDI	2.97
ASI	0.27
CEI-6G/SR/LR	6.375/4.976–6.375/4.976–6.375
CPRI	0.6144, 1.2288, 2.4576, 3.072
Ethernet-XAUI	4 x 3.1250
Fibre Channel	1.0625, 2.125, 4.25, 8.5
Gigabit Ethernet	1.25
GPON	1.244 uplink, 2.488 downlink
HiGig+	3.75
HiGig2	4.0625
HyperTransport 3.0	2.4, 2.8, 3.2
Interlaken	3.125–6.375
OBSAI	0.768, 1.536, 3.072
PCI Express 1.1/2.0	2.5, 5
PCI Express Cable	2.5
SONET OC-3/OC-12/OC-48	0.155, 0.622, 2.488
Serial RapidIO	1.25, 2.5, 3.125

Stratix FPGA series package and I/O matrices

368 Number indicates available user I/O pins.

Vertical migration (same V_{CC}, GND, ISP, and input pins). User I/Os may be less than labelled for vertical migration.

All Stratix series devices are offered in commercial and industrial temperatures and RoHS-compliant packages.

Stratix IV GX FPGAs (0.9V) up to 8.5-Gbps transceivers¹

Pin count		EP4SGX70D	EP4SGX110D	EP4SGX110F	EP4SGX230D	EP4SGX230F	EP4SGX230H	EP4SGX230K	EP4SGX290F	EP4SGX290H	EP4SGX290K	EP4SGX360F	EP4SGX360H	EP4SGX360K	EP4SGX530H	EP4SGX530K	EP4SGX530N
FBGA (F)	780 pin	368	368		368				288 ²			288 ²					
	1,152 pin			368		560			560			560					
	1,152 pin						560			560			560		560 ²		
	1,517 pin							736			736			736		736	
	1,932 pin																904

¹ All data is preliminary

² Hybrid package (flip chip) FBGA

Stratix IV E FPGAs (0.9V) High density, high performance, low power¹

480 Number indicates available user I/O pins.

Vertical migration (same V_{CC}, GND, ISP, and input pins).

All Stratix series devices are offered in commercial and industrial temperatures and RoHS-compliant packages.

Pin count		EP4SE110	EP4SE230	EP4SE290	EP4SE360	EP4SE530	EP4SE680
FBGA (F)	780 pin	480	480	480 ²	480 ²		
	1,152 pin			736	736	736 ²	736 ²
	1,517 pin			864	864	960	960
	1,760 pin					960	1,104

¹ All data is preliminary

² Hybrid package (flip chip) FBGA

Stratix FPGA series package and I/O matrices (continued)

342 Number indicates available user I/O pins.

Vertical migration (same V_{CC}, GND, ISP, and input pins).

All Stratix series devices are offered in commercial and industrial temperatures and RoHS-compliant packages.

		Stratix III L FPGAs Balanced logic, memory, and DSP						Stratix III E FPGAs Memory/DSP enhanced			
		EP3SL50	EP3SL70	EP3SL110	EP3SL150	EP3SL200	EP3SL340	EP3SE50	EP3SE80	EP3SE110	EP3SE260 ¹
FBGA (F)	484 pin (flip chip)	288	288					288			
	780 pin	480	480	480	480			480	480	480	
	1,152 pin			736	736				736	736	736
	1,517 pin					864	960				960
	780 pin						1,104				
HBGA (H)	780 pin										480
	1,152 pin					736	736				

¹ EP3SE260 FPGAs are the optimum solution for both logic (Stratix III L FPGA) and DSP/memory (Stratix III E FPGA) applications at this density.

342 Number indicates available user I/O pins.

Vertical migration (same V_{CC}, GND, ISP, and input pins). User I/Os may be less than labelled for vertical migration.

All Stratix series devices are offered in commercial and industrial temperatures and RoHS-compliant packages.

		Stratix II FPGAs (1.2V) High density, high performance						Stratix II GX FPGAs (1.2V) 6.375-Gbps transceivers							
		EP2S15	EP2S30	EP2S60	EP2S90	EP2S130	EP2S180	EP2SGX30C	EP2SGX30D	EP2SGX60C	EP2SGX60D	EP2SGX60E	EP2SGX90E	EP2SGX90F	EP2SGX130G
FBGA (F)	484 pin (flip chip)	342	342	334											
	672 pin (flip chip)	366	500	492											
	780 pin				534	534		361	361	364	364				
	1,020 pin			718	758	742	742								
	1,152 pin											534	558		
	1,508 pin				902	1,126	1,170							650	734
Hybrid FBGA (H)	484 pin				308										

342

Number indicates available user I/O pins.

Vertical migration (same V_{CC}, GND, ISP, and input pins). User I/Os may be less than labelled for vertical migration.

All Stratix series devices are offered in commercial and industrial temperatures and RoHS-compliant packages.

Stratix IV GX FPGAs (0.9V) up to 8.5-Gbps transceivers¹

		EP4SGX70D	EP4SGX110D	EP4SGX110F	EP4SGX230D	EP4SGX230F	EP4SGX230H	EP4SGX230K	EP4SGX290F
Density and speed	Equivalent LEs	72,600	105,600	105,600	228,000	228,000	228,000	228,000	291,200
	ALMs	29,040	42,240	42,240	91,200	91,200	91,200	91,200	116,480
	Registers ²	58,080	84,480	84,480	182,400	182,400	182,400	182,400	232,960
	M9K memory blocks	462	660	660	1,235	1,235	1,235	1,235	936
	M144K memory blocks	16	16	16	22	22	22	22	36
	Embedded memory (Mbits)	6.3	8.1	8.1	13.9	13.9	13.9	13.9	13.3
	MLAB memory (Kbits)	908	1,320	1,320	2,850	2,850	2,850	2,850	3,640
	Maximum 18-bit x 18-bit multipliers	384	512	512	1,288	1,288	1,288	1,288	832
	Speed grades (fastest to slowest)	-3, -4	-3, -4	-3, -4	-3, -4	-3, -4	-2, -3, -4	-2, -3, -4	-3, -4
Architectural features	Embedded processor available	Nios II processor							
	Global clock networks	16	16	16	16	16	16	16	16
	Regional clock networks	64	64	64	64	64	64	64	88
	Periphery clock networks	56	56	56	88	88	88	88	88
	PLLs/unique outputs	3/27	3/27	4/34	3/27	6/54	6/54	8/68	6/54
	Design security	✓	✓	✓	✓	✓	✓	✓	✓
	HardCopy series device support	✓	✓	✓	✓	✓	✓	✓	✓
	HardCopy companion device	HC4GX1	HC4GX2		HC4GX3				HC4GX4
I/O features	I/O voltage levels supported (V)	1.2, 1.5, 1.8, 2.5, 3.3 ³							
	I/O standards supported	LVTTTL, LVCMOS, PCI, PCI-X, LVDS, mini-LVDS, RSDS, LVPECL, Differential SSTL-15, Differential SSTL-18, Differential SSTL-2, Differential HSTL-12, Differential HSTL-15, Differential HSTL-18, SSTL-15 (I and II), SSTL-18 (I and II), SSTL-2 (I and II), 1.2V HSTL (I and II), 1.5V HSTL (I and II), 1.8V HSTL (I and II)							
	Medium-performance LVDS channels	128	128	128	128	192	192	192	192
	Number of LVDS channels, 1.6 Gbps (receive/transmit)	28/28	28/28	28/28	28/28	44/44	44/44	88/88	44/44
	Embedded DPA circuitry	✓	✓	✓	✓	✓	✓	✓	✓
	Series and differential OCT	✓	✓	✓	✓	✓	✓	✓	✓
	Programmable drive strength	✓	✓	✓	✓	✓	✓	✓	✓
	Transceiver (SERDES) data rate range	600 Mbps–8.5 Gbps with PCS + PMA							
	Transceiver (SERDES) channels	8	8	16	8	16	16	24	16
	Transceiver (SERDES) data rate range	600 Mbps–3.2 Gbps with PMA only							
	Transceiver (SERDES) channels ⁴	–	–	–	–	–	8	12	–
	PCI Express hard IP blocks	1	1	2	1	2	2	2	2
External memory interfaces	Memory devices supported	DDR3, DDR2, DDR, QDR II, RLDRAM II SDR							
Configuration file sizes	Configuration file size (Mbits)	52	52	52	102	102	102	102	140

¹ All data is preliminary

² This is the base core logic register count. The ALM can support three registers when used in LUTREG mode, which increases total register count by an additional 50 percent.

³ 3.3V compliant, requires a 3.0V power supply

⁴ These are additional transceivers, total transceiver count is sum of 8.5-Gbps transceivers plus 3.2-Gbps transceivers.

Stratix FPGA series features

Part 2 of 2

Stratix IV GX FPGAs (0.9V) up to 8.5-Gbps transceivers¹

		EP4SGX290H	EP4SGX290K	EP4SGX360F	EP4SGX360H	EP4SGX360K	EP4SGX530H	EP4SGX530K	EP4SGX530N
Density and speed	Equivalent LEs	291,200	291,200	353,600	353,600	353,600	531,200	531,200	531,200
	ALMs	116,480	116,480	141,440	141,440	141,440	212,480	212,480	212,480
	Registers ²	232,960	232,960	282,880	282,880	282,880	424,960	424,960	424,960
	M9K memory blocks	936	936	1,248	1,248	1,248	1,280	1,280	1,280
	M144K memory blocks	36	36	48	48	48	64	64	64
	Embedded memory (Mbits)	13.3	13.3	17.7	17.7	17.7	20.3	20.3	20.3
	MLAB memory (Kbits)	3,640	3,640	4,420	4,420	4,420	6,640	6,640	6,640
	Maximum 18-bit x 18-bit multipliers	832	832	1,040	1,040	1,040	1,024	1,024	1,024
	Speed grades (fastest to slowest)	-2, -3, -4	-2, -3, -4	-3, -4	-2, -3, -4	-2, -3, -4	-3, -4	-3, -4	-3, -4
Architectural features	Embedded processor available	Nios II processor							
	Global clock networks	16	16	16	16	16	16	16	16
	Regional clock networks	88	88	88	88	88	88	88	88
	Periphery clock networks	88	88	88	88	88	112	112	112
	PLLs/unique outputs	6/54	12/96	6/54	6/54	12/96	6/54	8/68	12/96
	Design security	✓	✓	✓	✓	✓	✓	✓	✓
	HardCopy series device support	✓	✓	✓	✓	✓	✓	✓	✓
	HardCopy companion device	HC4GX4		HC4GX5			HC4GX6		
I/O features	I/O voltage levels supported (V)	1.2, 1.5, 1.8, 2.5, 3.3 ³							
	I/O standards supported	LVTTTL, LVCMOS, PCI, PCI-X, LVDS, mini-LVDS, RSDS, LVPECL, Differential SSTL-15, Differential SSTL-18, Differential SSTL-2, Differential HSTL-12, Differential HSTL-15, Differential HSTL-18, SSTL-15 (I and II), SSTL-18 (I and II), SSTL-2 (I and II), 1.2V HSTL (I and II), 1.5V HSTL (I and II), 1.8V HSTL (I and II)							
	Medium-performance LVDS channels	192	192	192	192	192	192	192	256
	Number of LVDS channels, 1.6 Gbps (receive/transmit)	44/44	88/88	44/44	44/44	88/88	44/44	88/88	98/98
	Embedded DPA circuitry	✓	✓	✓	✓	✓	✓	✓	✓
	Series and differential OCT	✓	✓	✓	✓	✓	✓	✓	✓
	Programmable drive strength	✓	✓	✓	✓	✓	✓	✓	✓
	Transceiver (SERDES) data rate range	600 Mbps–8.5 Gbps with PCS + PMA							
	Transceiver (SERDES) channels	16	24	16	16	24	16	24	32
	Transceiver (SERDES) data rate range	600 Mbps–3.2 Gbps with PMA only							
	Transceiver (SERDES) channels ⁴	8	12	–	8	12	8	12	16
PCI Express hard IP blocks	2	2	2	2	2	4	4	4	
External memory interfaces	Memory devices supported	DDR3, DDR2, DDR, QDR II, RDRAM II SDR							
Configuration file sizes	Configuration file size (Mbits)	140	140	140	140	140	189	189	189

¹ All data is preliminary² This is the base core logic register count. The ALM can support three registers when used in LUTREG mode, which increases total register count by an additional 50 percent.³ 3.3V compliant, requires a 3.0V power supply⁴ These are additional transceivers, total transceiver count is sum of 8.5-Gbps transceivers plus 3.2-Gbps transceivers.

Stratix FPGA series features

Stratix IV E FPGAs (0.9V) High density, high performance, low power ¹

		EP4SE110	EP4SE230	EP4SE290	EP4SE360	EP4SE530	EP4SE680
Density and speed	Equivalent LEs	105,600	228,000	291,200	353,600	531,200	681,100
	ALMs	42,240	91,200	116,480	141,440	212,480	272,440
	Registers ²	84,480	182,400	232,960	282,880	424,960	544,880
	M9K memory blocks	660	1,235	936	1,248	1,280	1,529
	M144K memory blocks	16	22	36	48	64	64
	Embedded memory (Mbits)	8.1	13.9	13.3	17.7	20.3	22.4
	MLAB memory (Kbits)	1,320	2,850	3,640	4,420	6,640	8,514
	Maximum 18-bit x 18-bit multipliers	512	1,288	832	1,040	1,024	1,360
	Speed grades (fastest to slowest)	-2, -3, -4	-2, -3, -4	-2, -3, -4	-2, -3, -4	-3, -4	-3, -4
Architectural features	Embedded processor available	Nios II processor					
	Global clock networks	16	16	16	16	16	16
	Regional clock networks	88	88	88	88	88	88
	Periphery clock networks	56	88	88	88	112	112
	PLLs/unique outputs	6/54	12/96	6/54	6/54	12/96	6/54
	Design security	✓	✓	✓	✓	✓	✓
	HardCopy series device support	✓	✓	✓	✓	✓	✓
	HardCopy companion device	HC4GX1	HC4GX2	HC44	HC4GX3	HC46	HC47
I/O features	I/O voltage levels supported (V)	1.2, 1.5, 1.8, 2.5, 3.3 ³					
	I/O standards supported	LVTTTL, LVCMOS, PCI, PCI-X, LVDS, mini-LVDS, RSDS, LVPECL, Differential SSTL-15, Differential SSTL-18, Differential SSTL-2, Differential HSTL-12, Differential HSTL-15, Differential HSTL-18, SSTL-15 (I and II), SSTL-18 (I and II), SSTL-2 (I and II), 1.2V HSTL (I and II), 1.5V HSTL (I and II), 1.8V HSTL (I and II)					
	Medium-performance LVDS channels	128	128	256	256	256	288
	Number of LVDS channels, 1.6 Gbps (receive/transmit)	56/56	56/56	88/88	88/88	112/112	132/132
	Embedded DPA circuitry	✓	✓	✓	✓	✓	✓
	Series and differential OCT	✓	✓	✓	✓	✓	✓
	Programmable drive strength	✓	✓	✓	✓	✓	✓
	Memory devices supported	DDR3, DDR2, DDR, QDR II, RDRAM II SDR					
Configuration file sizes	Configuration file size (Mbits)	52	102	140	140	189	230

¹ All data is preliminary

² This is the base core logic register count. The ALM can support three registers when used in LUTREG mode, which increases total register count by an additional 50 percent.

³ 3.3V compliant, requires a 3.0V power supply

Stratix FPGA series features (continued)

		Stratix III L FPGAs Balanced logic, memory, and DSP						Stratix III E FPGAs Memory/DSP enhanced			
		EP3SL50	EP3SL70	EP3SL110	EP3SL150	EP3SL200	EP3SL340	EP3SE50	EP3SE80	EP3SE110	EP3SE260 ¹
Density and speed	Equivalent LEs	47,500	67,500	106,500	142,000	198,900	338,000	47,500	80,000	106,500	254,400
	ALMs	19,000	27,000	42,600	56,800	79,560	135,200	19,000	32,000	42,600	101,760
	Registers ²	38,000	54,000	85,200	113,600	159,120	270,400	38,000	64,000	85,200	203,520
	M9K memory blocks	108	150	275	355	468	1,040	400	495	639	864
	M144K memory blocks	6	6	12	16	24	48	12	12	16	48
	Embedded memory (Kbits)	1,836	2,214	4,203	5,499	7,668	16,272	5,328	6,183	8,055	14,688
	MLAB memory (Kbits)	594	844	1,344	1,781	2,500	4,219	594	1,000	1,344	3,188
	Maximum 18-bit x 18-bit multipliers	216	288	288	384	576	576	384	672	896	768
	Speed grades (fastest to slowest)	-2, -3, -4	-2, -3, -4	-2, -3, -4	-2, -3, -4	-2, -3, -4	-2, -3, -4	-2, -3, -4	-2, -3, -4	-2, -3, -4	-2, -3, -4
Architectural features	Embedded processor available	Nios II processor									
	Global clock networks	16	16	16	16	16	16	16	16	16	16
	Regional clock networks	48	48	48	48	88	88	48	48	48	88
	Periphery clock networks	104	104	208	208	208	208	104	208	208	208
	PLLs/unique outputs	4/34	4/34	8/68	8/68	12/96	12/96	8/68	12/96	12/96	12/96
	Design security	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	HardCopy series device support ³	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
I/O features	I/O voltage levels supported (V)	1.2, 1.5, 1.8, 2.5, 3.0									
	I/O standards supported	LVDS, LVPECL, Differential SSTL-18, Differential SSTL-2, Differential HSTL, SSTL-18 (I and II), SSTL-15 (I and II), SSTL-2 (I and II), 1.5V HSTL (I and II), 1.8V HSTL (I and II), PCI, PCI-X 1.0, LVTTTL, LVCMOS									
	True-LVDS maximum data rate (Mbps)	125–1,250									
	Number of LVDS channels (receive/transmit)	56/56	56/56	88/88	88/88	112/112	132/132	56/56	88/88	88/88	112/112
	Embedded DPA circuitry	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	Series and differential OCT	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	Programmable drive strength	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
External memory interfaces	Memory devices supported	DDR3, DDR2, DDR, QDR II, RLDRAM II, SDR									
Configuration file sizes	Configuration file size (Mbits) ³	22	22	47	47	66	120	26	48	48	93

¹ EP3SE260 FPGAs are the optimum solution for both logic (Stratix III L FPGA) and DSP/memory (Stratix III E FPGA) applications at this logic density.

² This is the base core logic register count. The ALM can support three registers when used in LUTREG mode, which can increase total register count by an additional 50 percent.

³ This data is preliminary.

Stratix FPGA series features (continued)

		Stratix II GX FPGAs (1.2V) 6.375-Gbps transceivers							
		EP2SGX30C	EP2SGX30D	EP2SGX60C	EP2SGX60D	EP2SGX60E	EP2SGX90E	EP2SGX90F	EP2SGX130G
Density and speed	Equivalent LEs	33,880	33,880	60,440	60,440	60,440	90,960	90,960	132,540
	ALMs	13,552	13,552	24,176	24,176	24,176	36,384	36,384	53,016
	ALUTs	27,104	27,104	48,352	48,352	48,352	72,768	72,768	106,032
	Total RAM (Kbits) ¹	1,338	1,338	2,485	2,485	2,485	4,415	4,415	6,590
	M512 RAM blocks (512 bits + 64 parity bits)	202	202	329	329	329	488	488	699
	M4K RAM blocks (4 Kbits ¹ + 512 parity bits)	144	144	255	255	255	408	408	609
	M-RAM blocks (512 Kbits ¹ + 65,536 parity bits)	1	1	2	2	2	4	4	6
	18-bit x 18-bit/9-bit x 9-bit embedded multipliers	64/128	64/128	144/288	144/288	144/288	192/384	192/384	252/504
	Speed grades (fastest to slowest)	-3, -4, -5	-3, -4, -5	-3, -4, -5	-3, -4, -5	-3, -4, -5	-3, -4, -5	-3, -4, -5	-3, -4, -5
Architectural features	Embedded processor available	Nios II processor							
	DSP blocks	16	16	36	36	36	48	48	63
	I/O registers per I/O element	6	6	6	6	6	6	6	6
	True dual-port RAM	✓	✓	✓	✓	✓	✓	✓	✓
	Global and regional clock networks	48	48	48	48	48	48	48	48
	PLLs/unique outputs	4/18	4/18	4/18	4/18	8/36	8/36	8/36	8/36
	Design security	✓	✓	✓	✓	✓	✓	✓	✓
	HardCopy II device support	—	—	—	—	—	—	—	—
I/O features	I/O voltage levels supported (V)	1.5, 1.8, 2.5, 3.3							
	I/O standards supported	LVDS, LVPECL, HyperTransport, Differential SSTL-18, Differential SSTL-2, Differential HSTL, SSTL-18 (I and II), SSTL-2 (I and II), 1.5V HSTL (I and II), 1.8V HSTL (I and II), PCI, PCI-X 1.0, LVTTTL, LVCMOS							
	True-LVDS maximum data rate (Mbps)	125–1,000							
	True-LVDS channels (receive/transmit)	31/29	31/29	31/29	31/29	42/42	47/45	59/59	73/71
	Embedded DPA circuitry	✓	✓	✓	✓	✓	✓	✓	✓
	Series and differential OCT	✓	✓	✓	✓	✓	✓	✓	✓
	Programmable drive strength	✓	✓	✓	✓	✓	✓	✓	✓
	Transceiver (SERDES) data rate range	600 Mbps–6.375 Gbps							
External memory interfaces	Transceiver (SERDES) channels	4	8	4	8	12	12	16	20
	Memory devices supported	DDR2, DDR, QDR II, RDRAM II, SDR							
Configuration file sizes	Configuration file size (Mbits)	10	10	17	17	17	28	28	40

¹ Kbits = 1,024 bits

Stratix FPGA series features (continued)

		Stratix II FPGAs (1.2V) High density, high performance					
		EP2S15	EP2S30	EP2S60	EP2S90	EP2S130	EP2S180
Density and speed	Equivalent LEs	15,600	33,880	60,440	90,960	132,540	179,400
	ALMs	6,240	13,552	24,176	36,384	53,016	71,760
	ALUTs	12,480	27,104	48,352	72,768	106,032	143,520
	Total RAM (Kbits) ¹	410	1,338	2,485	4,415	6,590	9,163
	M512 RAM blocks (512 bits + 64 parity bits)	104	202	329	488	699	930
	M4K RAM blocks (4 Kbits ¹ + 512 parity bits)	78	144	255	408	609	768
	M-RAM blocks (512 Kbits ¹ + 65,536 parity bits)	0	1	2	4	6	9
	18-bit x 18-bit/9-bit x 9-bit embedded multipliers	48/96	64/128	144/288	192/384	252/504	384/768
	Speed grades (fastest to slowest)	-3, -4, -5	-3, -4, -5	-3, -4, -5	-3, -4, -5	-3, -4, -5	-3, -4, -5
Architectural features	Embedded processor available	Nios II processor					
	DSP blocks	12	16	36	48	63	96
	I/O registers per I/O element	6	6	6	6	6	6
	True dual-port RAM	✓	✓	✓	✓	✓	✓
	Global and regional clock networks	48	48	48	48	48	48
	PLLs/unique outputs	6/28	6/28	12/56	12/56	12/56	12/56
	Design security	✓	✓	✓	✓	✓	✓
	HardCopy II device support	–	✓	✓	✓	✓	✓
I/O features	I/O voltage levels supported (V)	1.5, 1.8, 2.5, 3.3					
	I/O standards supported	LVDS, LVPECL, HyperTransport, Differential SSTL-18, Differential SSTL-2, Differential HSTL, SSTL-18 (I and II), SSTL-2 (I and II), 1.5V HSTL (I and II), 1.8V HSTL (I and II), PCI, PCI-X 1.0, LVTTTL, LVCMOS					
	True-LVDS maximum data rate (Mbps)	125–1,000					
	True-LVDS channels (receive/transmit)	38/38	58/58	80/84	114/118	152/156	152/156
	Embedded DPA circuitry	✓	✓	✓	✓	✓	✓
	Series and differential OCT	✓	✓	✓	✓	✓	✓
	Programmable drive strength	✓	✓	✓	✓	✓	✓
External memory interfaces	Memory devices supported	DDR2, DDR, QDR II, RDRAM II, SDR					
Configuration file sizes	Configuration file size (Mbits)	5	10	17	28	40	53

¹ Kbits = 1,024 bits

FOR MORE INFORMATION

Stratix IV FPGAs www.altera.com/stratix4
 Stratix III FPGAs www.altera.com/stratix3
 Stratix II FPGAs www.altera.com/stratix2
 Training www.altera.com/training

HardCopy ASIC series


Altera's HardCopy® system development methodology delivers the lowest total cost, lowest risk system designs. This methodology speeds the design of deep sub-micron HardCopy ASICs by employing the ease and flexibility of FPGA-based system design. This unique flow allows true hardware/software co-design and co-verification, so you can deliver your system to market sooner—on average 9 to 12 months faster than competing solutions.

Prototyping based on Stratix® series FPGAs allows you to get your system and system software/firmware ready prior to HardCopy ASIC design handoff. One design—using one register transfer level (RTL), one set of IP cores, and one tool (Quartus® II design software)—delivers both FPGA and ASIC implementations. Since all test insertion is managed by the HardCopy Design Center, you and your team won't need to spend any time on design for test. In addition, you'll spend zero time on design for manufacturability or design for yield. Altera's extensive design experience and solid, long-standing partnership with TSMC means excellent built-in manufacturability, yield, and reliability.

Key features

- System development methodology
 - Design-once seamless prototyping
 - True software/hardware co-design and co-verification
 - 9- to 12-month faster system availability
 - Fast and predictable design handoff to production HardCopy ASIC silicon
- Dramatically lower power, increased SEU immunity, increased security, and lower bill of materials (BOM) cost compared to FPGAs
- Lower embedded systems costs with Nios® II embedded processors

HardCopy ASIC



HardCopy® IV

40-nm HardCopy ASICs with 6.5+-Gbps transceiver option

The HardCopy IV ASIC family offers the benefits of FPGAs and the benefits of ASICs. HardCopy IV ASICs deliver the lowest risk, lowest total cost, and fastest time to market available in leading-edge ASIC technology today, as well as risk-free, seamless prototyping on a Stratix IV family FPGA. The HardCopy IV family meets your diverse application needs in markets including wireless, wireline, high-performance computing, high-reliability computing, storage, military, and more. The HardCopy IV family has two family variants: HardCopy IV GX and HardCopy IV E devices. Both the HardCopy IV GX and HardCopy IV E variants include high-performance logic—up to 11.5M and 13.3M ASIC gates respectively, and up to 16.8M bits of memory—as well as DSP capabilities. Additionally, HardCopy IV GX ASICs feature up to 24 transceiver channels with up to 6.5+-Gbps bandwidth.

HardCopy IV family features summary (GX only)

Equivalent SERDES block to Stratix IV GX FPGAs	Same transceivers, same hard IP as Stratix IV GX companion devices.
Multi-gigabit transceivers (PMA and PCS features)	Up to 24 high-speed transceiver channels (full duplex) with PLL-based CDR operating from 600 Mbps to 6.5 Gbps (6.5 Gbps industrial) and superior signal integrity.
Enhanced transceiver block	Two central CMUs per transceiver block, used for reference clock inputs and transmit PLLs. CMUs can be configured as transmit PLLs or as additional channels 5 and 6 in transceiver block. Embedded PCS pattern detector, word aligner with programmable pattern byte reordering, 8b/10b encoder/decoder, protocol-specific hardware for PCI Express (PIPE interface), XAUI, Gigabit Ethernet, and SONET.
Protocol support	Support for PCI Express Gen1 and Gen2, XAUI/HiGig-XAUI, Gigabit Ethernet, Serial RapidIO®, Interlaken, CEI-6G, SONET OC-3/12/48, SDI, ASI, GPON, SATA, 3G and 6G basic modes with rule-based protocol configuration.
Plug & Play Signal Integrity	Support for adaptive equalization and hot socketing.
PCI Express hard IP	Complete protocol solution with embedded PCI Express hard IP blocks, v1.1 and v2.0, that implement PHY-MAC layer, data-link layer, and transaction layer functionality.
Dynamic transceiver reconfiguration	Provides the capability to support multiple protocols and data rates on the same transceiver channel in a single HardCopy ASIC.
Low-power transceiver	Typical PMA power consumption of 100 mW at 3.125 Gbps and 135 mW at 6.375 Gbps per channel.

HardCopy IV family features summary (E and GX)

Seamless prototyping	HardCopy IV ASIC seamless prototyping on Stratix IV FPGAs is available for both GX and E variants. Single toolset across HardCopy IV and Stratix IV devices for enhanced productivity. All of the HardCopy IV and Stratix IV devices are supported in Quartus II software.
External memory interface circuitry	HardCopy IV ASICs support external DRAM and SRAM memory interfaces, including DDR3, DDR2, QDR II, QDR, RLD RAM II, and RLD RAM. The dynamic self-calibrating PHY continuously compensates for PVT variations during operation. You can achieve higher performance, up to 800 Mbps (400 MHz), better signal integrity, increased robustness, expanded design margin, and lower power in your designs. Complete solutions including easy-to-use IP, hardware development platforms, and comprehensive timing analysis are available to minimize risk and speed time to market.
TriMatrix memory	TriMatrix memory offers three memory block sizes to suit complex design needs: 640-bit MLABs, 9-Kbit M9K blocks, and 144-Kbit M144K blocks. With up to 16.8 Mbits of memory performing at over 600 MHz and features such as dual-port RAM and integrated ECC, TriMatrix memory is more flexible and efficient, and provides higher memory bandwidth than any other memory architecture. All three configurations support parity checking.
DSP blocks	HardCopy IV ASICs support up to 1,288 multipliers. All DSP block circuitry used will be implemented in HCell logic, the basic building block used for logic implementation. DSP blocks include all associated pipelining, adders, and accumulators. They are configurable to support 9-bit x 9-bit, 12-bit x 12-bit, 18-bit x 18-bit, or 36-bit x 36-bit multipliers at up to 455 MHz and 586 GMACS, making HardCopy IV ASICs ideal for custom-built DSP replacement, as well as coprocessing for hardware acceleration.
High-speed single-ended I/O support	Up to 20 modular I/O banks with programmable slew rate, programmable drive strength, programmable output delay, OCT, and dynamic trace compensation. Maximizes I/O flexibility and performance for over 40 industry standards to optimally match your system requirements. I/O configuration will be fixed after design handoff to design center.
High-speed differential I/O support	Optimized LVDS I/O pins at between 150-MHz and 1.25-GHz performance provide high performance and excellent signal integrity with half the capacitance of competitive devices. HardCopy IV LVDS supports hard DPA and SERDES blocks with clock forwarding for CDR in high-speed chip-to-chip applications. MegaWizard® Plug-Ins make differential I/O configuration straightforward and easy.
DPA	As a feature of the LVDS high-speed differential I/O support, DPA minimizes bit errors and simplifies PCB layout and timing management for high-speed data transfer. DPA eliminates channel-to-channel and channel-to-clock skew in high-speed data transmission systems.
OCT	Serial, parallel, dynamic, and differential OCT simplifies design and reduces component count. Dynamic OCT enables support for next-generation memory interface designs that must dynamically switch between serial and parallel termination. Digital calibration circuitry provides industry-leading OCT tolerance within 10 percent.
Clock management features	Up to 12 PLLs with up to 10 unique customizable outputs per PLL ranging from 5 to 720 MHz. There are also up to 16 global, 88 quadrant, and 208 periphery clocks. Advanced clock management features include PLL reconfiguration, frequency synthesis, programmable phase shift, programmable delay shift, external feedback, and programmable bandwidth for all on- and off-chip clocking requirements.
Design security	Hardwired logic is very secure—no bitstream, no external programming device.
Hot socketing and power sequencing	Robust on-chip hot-socketing and power-sequencing support ensures proper device operation independent of the power-up sequence.
Increased SEU immunity	The connections between HCells are hardwired after via programming. The high SEU tolerance nature of HardCopy IV ASICs is due not only to the hardwiring, but also to improved architecture of sequential elements.
Nios® II embedded processor support	Nios II embedded processors reduce cost, increase flexibility, and offer an ideal replacement for low-cost discrete microprocessors. Nios II soft processor support for HardCopy IV ASICs offers over 300-DMIPS performance.

HardCopy IV transceiver protocol support (GX only)

Protocol	Data rate
3G SDI	2.97
ASI	0.27
CEI-6G/SR/LR	6.375/4.976–6.375/4.976–6.375
CPRI	0.6144, 1.2288, 2.4576, 3.072
DDR-XAUI	6.25
Ethernet-XAUI	4 x 3.125
Fibre Channel	1.0625, 2.125, 4.25, 6.52
Gigabit Ethernet	1.25
GPON	1.244 uplink, 2.488 downlink
HiGig+	3.75
HiGig2	4.0625
Interlaken	3.125–6.375
OBSAI	0.768, 1.536, 3.072
PCI Express Gen1, Gen2	2.5, 5.0
PCI Express Cable	2.5
SAS	1.5, 3.0, 6.0
SATA	1.5, 3.0
SDI SD/HD	0.27/1.485
SerialLite II	0.6–6.375
Serial RapidIO	1.25, 2.5, 3.125
SFI5.1	2.488–3.125 (up to 8 channels)
SONET OC-3/OC-12/OC-48	0.155, 0.622, 2.488
SPAUI	3.2, 6.4



HardCopy ASICs deliver high levels of integration

HardCopy III ASICs deliver low risk, low total cost, and rapid time to market and profit for your custom logic needs. The HardCopy methodology allows you to prototype your system with Stratix III FPGAs and completely prepare your system for production. Altera's HardCopy Design Center uses a fast and predictable turnkey process to implement the low-cost, low-power, functionally equivalent, pin-compatible HardCopy III device. Our methodology is the low-risk solution for rapid system development and production.

HardCopy III family features summary

Seamless prototyping	HardCopy III ASIC seamless prototyping on Stratix III FPGAs. Single toolset across HardCopy III and Stratix III devices for enhanced productivity. All of the HardCopy III and Stratix III devices supported in Quartus II software.
External memory interface circuitry	HardCopy III ASICs support external DRAM and SRAM memory interfaces, including DDR3, DDR2, QDR II, QDR, RDRAM II, and RDRAM. The dynamic self-calibrating PHY continuously compensates for PVT variations during operation. You can achieve higher performance, up to 800 Mbps (400 MHz), better signal integrity, increased robustness, expanded design margin, and lower power in your designs. Complete solutions including easy-to-use IP, hardware development platforms, and comprehensive timing analysis are available to minimize risk and speed time to market.
TriMatrix memory	TriMatrix memory offers three memory block sizes to suit complex design needs: 640-bit MLABs, 9-Kbit M9K blocks, and 144-Kbit M144K blocks. With up to 16.3 Mbits of memory performing at over 600 MHz and features such as dual-port RAM and integrated ECC, TriMatrix memory is more flexible and efficient, and provides higher memory bandwidth than any other memory architecture. All three configurations support parity checking.
DSP blocks	HardCopy III ASICs support up to 896 multipliers. All used DSP block circuitry is implemented in HCell logic. DSP blocks include all associated pipelining, adders, and accumulators. They are configurable to support 9-bit x 9-bit, 12-bit x 12-bit, 18-bit x 18-bit, or 36-bit x 36-bit multipliers at up to 455 MHz and 407 GMACS, making HardCopy III ASICs ideal for custom-built DSP replacement, as well as co-processing for hardware acceleration.
High-speed single-ended I/O support	Up to 20 modular I/O banks with programmable slew rate, programmable drive strength, programmable output delay, OCT, and dynamic trace compensation. Maximizes I/O flexibility and performance for over 40 industry standards to optimally match your system requirements. I/O configuration will be fixed after Design Review 2.
High-speed differential I/O support	Optimized LVDS I/O pins at 150-MHz and 1.25-GHz performance provide high performance and excellent signal integrity with half the capacitance of competitive devices. HardCopy III LVDS supports hard DPA and SERDES blocks with clock forwarding for CDR in high-speed chip-to-chip applications. MegaWizard Plug-Ins make differential I/O configuration straightforward and easy.
DPA	As a feature of the LVDS high-speed differential I/O support, DPA minimizes bit errors and simplifies PCB layout and timing management for high-speed data transfer. DPA eliminates channel-to-channel and channel-to-clock skew in high-speed data transmission systems.
OCT	Serial, parallel, dynamic, and differential OCT simplifies design and reduces component count. Dynamic OCT enables support for next-generation memory interface designs that must dynamically switch between serial and parallel termination. Digital calibration circuitry provides industry-leading OCT tolerance within 10 percent.
Clock management features	Up to 12 PLLs with up to 10 unique customizable outputs per PLL ranging from 5 to 720 MHz. There are also up to 16 global, 88 quadrant, and 208 periphery clocks. Advanced clock management features include PLL reconfiguration, frequency synthesis, programmable phase shift, programmable delay shift, external feedback, and programmable bandwidth for all on- and off-chip clocking requirements.
Design security	Hardwired logic is very secure—no bitstream and no external programming device.
Hot socketing and power sequencing	Robust on-chip hot-socketing and power-sequencing support ensures proper device operation independent of the power-up sequence.
Increased SEU immunity	The connections between HCells are hardwired after via programming. The high SEU tolerance nature of the HardCopy III ASICs is due not only to the hardwiring, but also to improved architecture of sequential elements. HCells are the logic building blocks for HardCopy ASICs.
Nios II embedded processor support	Nios II embedded processors reduce cost, increase flexibility, and offer an ideal replacement for low-cost discrete microprocessors. Nios II soft processor support for HardCopy III ASICs offers over 300-DMIPS performance.



90-nm HardCopy ASICs with guaranteed success

The 1.2V, 90-nm HardCopy II family offers up to 3.6M ASIC gates, 8.8M bits of memory, and over 350-MHz system performance. Compared to FPGAs, HardCopy II ASICs offer over 50 percent power reduction and up to 100 percent higher core performance. For data on increased SEU immunity, contact your Altera representative. Using Stratix II FPGAs for seamless prototyping, plus our fast and predictable back-end process, provides fast time to market with reduced risk for your HardCopy ASIC-based systems.

HardCopy II family features summary

Clock management	Up to 12 programmable PLLs, providing robust clock management and frequency synthesis capabilities for maximum system performance. The PLLs provide high-end features, including clock switchover, PLL reconfiguration, spread-spectrum clocking, frequency synthesis, programmable phase shift, programmable delay shift, external feedback, and programmable bandwidth.
Differential I/O support	Support for high-speed differential I/Os for data rates up to 1 Gbps to address the high-performance needs of emerging I/O interfaces, including support for the LVDS, LVPECL, and HyperTransport standards.
External memory interfaces	Advanced external memory interface support allows you to integrate external high-density SRAM and DRAM devices into complex system designs without degrading data-access performance.
OCT	Series and differential OCT support can simplify board layout by minimizing the number of external resistors needed.
Single-ended I/O standards	High-bandwidth, single-ended I/O interface standards (SSTL, HSTL, PCI, and PCI Express) supported.
Source-synchronous protocols	Support for a wide array of high-speed interface standards, including SPI-4.2, SFI-4, 10 Gigabit Ethernet XSBI, and Serial RapidIO.
TriMatrix memory	Up to 8.8 Mbits of RAM—configurable M4K and M-RAM embedded RAM blocks for a wide range of features.

HardCopy ASIC series package and I/O matrix

HardCopy IV ASICs (0.9V) with 6.5+-Gbps transceivers option¹

All HardCopy series devices are offered in commercial and industrial temperatures and RoHS-compliant packages.

		HC4GX1YZ ⁵	HC4GX2YZ	HC4GX3YZ	HC4GX4YZ	HC4GX5YZ	HC4GX6YZ	HC4E2YZ	HC4E3YZ	HC4E4YZ	HC4E5YZ	HC4E6YZ	HC4E7YZ
FBGA (F)	484 pin (WF ²)							✓	✓				
	484 pin (LF ³)	✓	✓	✓				✓	✓				
	780 pin (WF)							✓	✓	✓			
	780 pin (LF)	✓	✓	✓	✓	✓		✓	✓	✓	✓		
	1,152 pin (LF)		✓	✓	✓	✓				✓	✓		
	1,152 pin (FF ⁴)			✓	✓	✓	✓			✓	✓	✓	✓
	1,517 pin (FF)			✓	✓	✓	✓			✓	✓	✓	✓

¹ All data is preliminary

² WF = wire bond

³ LF = Cost-optimized flip chip

⁴ FF = Performance-optimized flip chip

⁵ Y = I/O count, Z = package type

HardCopy III ASICs (0.9V)¹

All HardCopy series devices are offered in commercial and industrial temperatures and RoHS-compliant packages.

		HC31YZ	HC32YZ	HC33YZ	HC35YZ	HC36YZ	HC37YZ
FBGA (F)	484 pin (WF ²)	✓	✓	✓	✓	✓	✓
	484 pin (FF ³)	✓	✓	✓	✓	✓	✓
	780 pin (WF)	✓	✓	✓	✓	✓	✓
	780 pin (FF)	✓	✓	✓	✓	✓	✓
	1,152 pin (LF ⁴)		✓	✓	✓	✓	✓
	1,152 pin (FF)		✓	✓	✓	✓	✓
	1,517 pin (LF)				✓	✓	✓
	1,517 pin (FF)				✓	✓	✓

¹ All data is preliminary

² WF = wire bond

³ FF = Performance-optimized flip chip

⁴ LF = Cost-optimized flip chip

HardCopy II ASICs (1.2V)

951 Number indicates available user I/O pins.

All HardCopy series devices are offered in commercial and industrial temperatures and RoHS-compliant packages.

		HC210W	HC210	HC220W	HC220	HC230	HC240
FBGA (F)	484 pin (wire bond)	308					
	484 pin (flip chip)		334				
	672 pin (wire bond)			440			
	672 pin (flip chip)				492		
	780 pin (wire bond)			445			
	780 pin				494		
	1,020 pin					698	742
	1,508 pin						951

HardCopy ASIC series features

HardCopy IV ASICs (0.9V) with 6.5+-Gbps transceivers option¹

		HC4GX1YZ	HC4GX2ZY	HC4GX3YZ	HC4GX4YZ	HC4GX5YZ	HC4GX6YZ	HC4E2YZ	HC4E3YZ	HC4E4YZ	HC4E5YZ	HC4E6YZ	HC4E7YZ
Density and speed	Usable ASIC gates	2.8M	3.9M	9.2M	7.6M	9.5M	11.5M	3.9M	9.2M	7.6M	9.5M	11.5M	13.3M
	Embedded memory (Mbits)	6.3	8.1	9.8–12.2	10.6–12.7	10.6–13.3	13.3	8.1	10.7	12.1–13.3	16.8	16.8	16.8
	Maximum 18-bit x 18-bit multipliers	384	512	1,288	832	1,040	1,024	512	1,288	832	1,040	1,024	1,024
Architectural features	Embedded processor available	Nios II processor											
	PLLs	3	3/4	3/6/8	3/6/8	3/6/8	6/8	4	4	4/8/12	4/8/12	8/12	8/12
	Design security ²	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	Stratix series prototyping support	EP4SGX 70	EP4SGX 110	EP4SGX 230	EP4SGX 290	EP4SGX 360	EP4SGX 530	EP4SE 110	EP4SE 230	EP4SE 290	EP4SE 360	EP4SE 530	EP4SE 680
I/O features	I/O voltage levels supported	1.2, 1.5, 1.8, 2.5, 3.0											
	I/O standards supported	LVTTTL, LVCMOS, PCI, PCI-X, LVDS, mini-LVDS, RSDS, LVPECL, Differential SSTL-15, Differential SSTL-18, Differential SSTL-2, Differential HSTL-12, Differential HSTL-15, Differential HSTL-18, SSTL-15 (I and II), SSTL-18 (I and II), SSTL-2 (I and II), 1.2V HSTL (I and II), 1.5V HSTL (I and II), 1.8V HSTL (I and II)											
	True-LVDS maximum data rate (Mbps)	150–1,250+											
	Embedded DPA circuitry	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	Series and differential OCT	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	Programmable drive strength ³	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
	Transceiver (SERDES) data rate range	600 Mbps–6.5+ Gbps with PCS + PMA ⁴											
	6.5-Gbps transceiver (SERDES) channels	4/8	4/8/16	4/8/16/24	16/24	16/24	24	–	–	–	–	–	–
External memory interfaces	Memory devices supported	DDR3, DDR2, DDR, QDR II, RLDRAM II SDR											

¹ All data is preliminary

² Since all HardCopy ASICs contain hardwired logic, they are inherently secure.

³ Settings are fixed post Design Review 2.

⁴ Subject to increase, pending characterization.

HardCopy ASIC series features

		HardCopy III ASICs (0.9V) ¹					
		HC31YZ	HC32YZ	HC33YZ	HC35YZ	HC36YZ	HC37YZ
Density and speed	Usable ASIC gates	2.7M	3.6M	5.8M	5.3M	5.3M	7.0M
	Embedded memory (Mbits)	4.2	5.5	5.5–8.0	5.5–9.6	5.5–14.7	11.3–16.3
	Maximum 18-bit x 18-bit multipliers	288	384	896	576	768	576
Architectural features	Embedded processor available	Nios II processor					
	DSP blocks	Implemented in HCell macros					
	True dual-port RAM	✓	✓	✓	✓	✓	✓
	PLLs	4	4/8	4/8	4/8/12	4/8/12	4/8/12
	Design security ²	✓	✓	✓	✓	✓	✓
	Stratix series prototyping support	EP3SL110	EP3SL150	EP3SE110	EP3SL200	EP3SE260	EP3SL340
I/O features	I/O voltage levels supported	1.2, 1.5, 1.8, 2.5, 3.0					
	I/O standards supported	LVTTTL, LVCMOS, PCI, PCI-X, LVDS, mini-LVDS, RSDS, LVPECL, Differential SSTL-15, Differential SSTL-18, Differential SSTL-2, Differential HSTL-12, Differential HSTL-15, Differential HSTL-18, SSTL-15 (I and II), SSTL-18 (I and II), SSTL-2 (I and II), 1.2V HSTL (I and II), 1.5V HSTL (I and II), 1.8V HSTL (I and II)					
	True-LVDS maximum data rate (Mbps)	150–1,250+					
	Embedded DPA circuitry	✓	✓	✓	✓	✓	✓
	Series and differential OCT	✓	✓	✓	✓	✓	✓
	Programmable drive strength ³	✓	✓	✓	✓	✓	✓
External memory interfaces	Memory devices supported	DDR3, DDR2, DDR, QDR II, RDRAM II SDR					

¹ All data is preliminary.

² Since all HardCopy ASICs contain hardwired logic, they are inherently secure.

³ Settings are fixed post Design Review 2.

HardCopy ASIC series features

		HardCopy II ASICs (1.2V)					
		HC210W	HC210	HC220W	HC220	HC230	HC240
Density and speed	Usable ASIC gates	1.0M	1.0M	1.6M	1.6M	2.2M	2.2M
	Additional gates for DSP blocks	0	0	300,000	300,000	700,000	1,400,000
	LEs	–	–	–	–	–	–
	Total RAM bits	875,520	875,520	3,059,712	3,059,712	6,368,256	8,847,360
	M512 RAM blocks (512 bits + 64 parity bits)	–	–	–	–	–	–
	M4K RAM blocks (4 Kbits ¹ + 512 parity bits)	190	190	408	408	614	768
	M-RAM blocks (512 Kbits ¹ + 65,536 parity bits)	0	0	2	2	6	9
	Speed grades (fastest to slowest)	–	–	–	–	–	–
Architectural features	Embedded processor available	Nios II processor					
	DSP blocks	Implemented in HCell macros					
	18-bit x 18-bit/9-bit x 9-bit embedded multipliers	Implemented in HCell macros					
	I/O registers per I/O element	6	6	6	6	6	6
	True dual-port RAM	✓	✓	✓	✓	✓	✓
	Global and regional clock networks	16/32	16/32	16/32	16/32	16/32	16/32
	PLLs	4	4	4	4	8	12
I/O features	I/O voltage levels supported (V)	1.5, 1.8, 2.5, 3.3					
	I/O standards supported	LVDS, LVPECL, HyperTransport, Differential SSTL-18, Differential SSTL-2, Differential HSTL, SSTL-18 (I and II), SSTL-2 (I and II), 1.5V HSTL (I and II), 1.8V HSTL (I and II), PCI, PCI-X 1.0, LVTTTL, LVCMOS					
	External memory device interfaces	QDR II, DDR2, RLDRAM II, DDR, SDR					
	True-LVDS maximum data rate (Mbps)	125–1,000					
	True-LVDS channels (receive/transmit)	13/17	19/21	29/31	29/31	42/42	118/118
	Medium-speed LVDS data rate (Mbps) (receive/transmit)	–	–	–	–	–	–
	Series and differential OCT	✓	✓	✓	✓	✓	✓
	FPGA prototype options	EP2S30 EP2S60 EP2S90	EP2S30 EP2S60 EP2S90	EP2S30 EP2S60 EP2S90 EP2S130	EP2S60 EP2S90 EP2S130	EP2S90 EP2S130 EP2S180	EP2S180

¹ Kbits = 1,024 bits

FOR MORE INFORMATION

HardCopy ASIC series www.altera.com/hardcopy
 HardCopy IV ASICs www.altera.com/hardcopy4
 HardCopy III ASICs www.altera.com/hardcopy3
 HardCopy II ASICs www.altera.com/hardcopy2
 Training www.altera.com/training

Nios II embedded processors

When you're poised to capture an early-to-market advantage with your novel idea, turn to Altera's Nios® II embedded processor to give you a jump on the competition. By allowing you to drag and drop the precise mix of processors and peripherals needed to build your system, the versatile 32-bit Nios II processor equips you to create an exact-fit processor system in just minutes.

Compared with traditional embedded processors, the Nios II processor offers much more flexibility. It provides a customizable feature set, along with costs and performance that you can tailor to your unique system requirements. Used by each of the world's top 20 OEMs and with more than 20,000 development kits sold worldwide, the proven Nios II processor is the industry's most widely used soft processor as confirmed by the market research firm Gartner.

Nios® II Versatility

We designed Nios II processors to meet the unique demands of new design cycles, enabling you to:

- Choose the exact set of processors, peripherals, memory, and interfaces needed for your application
- Increase system performance without changing your board design or increasing clock frequency
- Eliminate the risk of processor and ASSP obsolescence
- Lower overall system cost, complexity, and power consumption with multiple functions combined on a single chip
- Remotely upgrade your designs in the field to improve competitiveness and address changing requirements
- Target any Stratix® or Cyclone® series FPGA or HardCopy® series ASIC
- Migrate designs to standard cell ASICs through the Synopsys DesignWare Star IP Program


Broad software support

The Nios II processor is backed by the Eclipse-based Nios II Integrated Development Environment (IDE) and a full range of software and operating system support provided by Altera and our partners. To boost performance, just right-click to accelerate your ANSI-C code using the Nios II C-to-Hardware Acceleration (C2H) Compiler (*see sidebar*).

Performance

With performance clocked at well over 300 DMIPS, the Nios II processor is the industry's fastest FPGA-based configurable processor. You'll uncover the true processing performance potential of your FPGAs when you offload their processing tasks from software into dedicated hardware accelerators. Unlike off-the-shelf processors, FPGAs have parallel processing capabilities that let you create very high-performance, high-throughput systems.

Nios II processor development flow



Nios II C2H Compiler

An easy-to-use productivity tool, the Nios II C2H Compiler lets you boost the performance of time-critical ANSI-C functions by converting them into hardware accelerators in the FPGA. With the Nios II C2H Compiler, you can:

- Accelerate performance of your Nios II embedded software—without increasing your clock frequency
- Design using the standard ANSI-C programming language and the familiar Eclipse-based Nios II IDE
- Boost software performance from 10X to 70X

Traditionally, offloading software to hardware accelerators has been a manual task, benefiting only those developers with the tools, experience, and time required to create, test, and integrate RTL blocks into their processor systems. The Nios II C2H Compiler automates the creation and integration of hardware accelerators, reducing development time from weeks to minutes.

Software and operating system support

	Product name	Vendor name	Description
Operating system	MicroC/OS-II	Micrium	Real-time kernel; included in Nios II development kits
	eCos	eCosCentric	Open-source real-time operating system
	Nucleus PLUS	Mentor Graphics	Royalty-free, real-time operating system
	ThreadX	Express Logic	High-performance, real-time kernel
	µClinux	SLS	Complete µClinux distribution for the Nios II processor
	Euros	Euros RTOS	Real-time operating system
	embOS	Segger	Royalty-free real-time operating system
	OSEK/VDX-OS	Vector	Preemptive multitasking operating system
Debuggers/ IDEs	Evidence	Erika Enterprise	Real-time operating system with microprocessor support
	Nios II IDE	Altera Corporation	Full-featured IDE included in the Nios II Embedded Design Suite (EDS)
	EDGE IDE	Mentor Graphics	Eclipse-based IDE and debugger
	TRACE32-PowerView	Lauterbach	Flexible, fast debug environment
	Universe	Adveda	Hardware/software co-verification tools
	System Navigator	First Silicon Solutions	Software debug probe
Compilers	TRACE32-PowerView	Lauterbach	Software debug probe environment
	Tasking VX-Toolset	Altium	Optimizing C compiler
	GCC Compiler	GNU	Standard GNU compiler for Nios II processor

Nios II performance features and summary

High-performance processor core	Optimize performance-critical applications with the Nios II/f “fast” core, which has 6-stage pipeline, dynamic branch prediction, instruction and data cache, and 300-DMIPS performance.
Multi-processor systems	Use multi-core systems to scale a system’s performance or to break up software applications into simpler tasks. The Nios II EDS includes support for creating customized multi-core systems using Nios II processors.
High-bandwidth bus structure	Automatically generate a system interconnect fabric to support any system that you build by using the SOPC Builder system generation tool, allowing you to generate high-throughput systems supporting simultaneous multiple master/slave connections, direct memory access (DMA) channels, and on-chip data buffers.
Hardware accelerators	Use logic and memory resources in the FPGA to offload and accelerate tasks that are typically implemented in application software. You can automate this process with the Nios II C2H Compiler.
Custom instructions	Accelerate time-critical software algorithms by adding custom instructions to the Nios II processor instruction set.
Fast configurable on-chip memory	Create fixed low-latency on-chip memory buffers for performance-critical applications.

Designing with Nios II processors

Nios II processors come with a comprehensive tool suite that raises productivity by helping you and your team efficiently build and manage customized embedded processor systems. What’s more, our software design suite was developed to accommodate the unique requirements of systems on programmable chips.

Nios II EDS

The Nios II EDS is a collection of all the tools, utilities, libraries, and drivers needed to develop embedded software for the Nios II processor. Among many other tools and utilities, the Nios II EDS includes the Nios II IDE, the primary software development environment for the Nios II processor. Within this Eclipse-based GUI, you can edit, compile, download, debug, and program flash devices.

Embedded partners

Working with Altera, you’ll have the backing of more than a dozen embedded partners, offering a broad range of third-party operating systems, middleware, and embedded software development tools.

Nios II EDS contents

- Code development tool: Nios II IDE
 - New project wizards
 - Software templates
 - Source navigator and editor
 - Compiler for C and C++ (GNU)
 - Instruction set simulator (ISS)
 - Based on the Eclipse community project
- Source debugger
- Flash programmer
 - Hardware Abstraction Layer (HAL)
 - MicroC/OS-II real-time operating system¹
 - NicheStack TCP/IP Network Stack—Nios II Edition¹
 - Newlib ANSI-C standard library
 - Simple file system
- Other Altera® command line tools and utilities
- Design examples
- C acceleration tool: Nios II C2H Compiler¹

¹ Production license sold separately

Hardware development tools

- Quartus II design software
- SOPC Builder, an exclusive Quartus II software tool that lets you build and evaluate systems at the block level easily and quickly
- SignalTap® II embedded logic analyzer plug-in for the Nios II processor
- FPGAView software from First Silicon Solutions for configuring and debugging Altera FPGA devices with Tektronix logic analyzers

Processor cores

The Nios II processor family consists of three CPU cores, each optimized for a specific price and performance range. All three CPU cores share a common 32-bit instruction set architecture, are binary code compatible, and are supported by the same software design suite. Simply choose which CPU is appropriate for each of your designs. You can also create multi-core systems with any Nios II CPU to scale a system's performance or to break up software applications into simpler tasks.

Nios II processor family members

Feature	Nios II/f (fast) processor	Nios II/s (standard) processor	Nios II/e (economy) processor
Description	Optimized for maximum performance	Balanced cost and performance	Optimized for minimum logic usage
Pipeline	6 stage	5 stage	1 stage
Multiplier	1 cycle ¹	3 cycle	Emulated in software
Branch prediction	Dynamic	Static	None
Instruction cache	Configurable	Configurable	None
Data cache	Configurable	None	None
Custom instructions	Up to 256	Up to 256	Up to 256

¹ Using DSP blocks in Stratix or Stratix II FPGAs

Embedded IP

SOPC Builder includes a portfolio of standard embedded peripherals that you can use to jump-start your system design. These IP blocks include HAL drivers that will let you immediately begin your software development. Visit www.altera.com/sopcready for the latest list of SOPC Builder-ready IP.

Altera embedded peripherals

• JTAG UART • UART • DMA • Timer • PIO • SPI • SDRAM controller • DDR SDRAM controller • Common flash interface controller	• Tri-state bridge • Serial flash controller (EPCS) • Ethernet • PLL core • Mailbox and mutex • LCD controller • System ID peripheral • Performance counter • SG-DMA
--	--

Ready to go?

Make a Nios II development kit your starting point—the kits include everything you'll need to begin designing embedded systems for FPGAs:

- Processor and peripheral IP (with perpetual use license)
- Complete software development tool suite
- FPGA development board (see the Development Kits section of this catalog for more details)

When you're ready to ship your product, you'll need the Nios II core license that's included with all Nios II development kits. You can also buy the license as a standalone offering. Contact your local Altera representative for details. This royalty-free license never expires and allows you to target your processor design to any Altera FPGA, so your software application is preserved even if the underlying hardware changes.

Start working with the Nios II processor today using free Altera development tools. Hardware engineers should download Quartus II Web Edition design software to create Nios II processor-based hardware systems. Software designers should download the Nios II EDS to begin programming Nios II applications right away.

FOR MORE INFORMATION

Nios II processors www.altera.com/nios
 Nios II development kits www.altera.com/nioskits
 Nios II C2H Compiler www.altera.com/c2h
 Nios user forum www.niosforum.org
 Nios wiki nioswiki.jot.com


Intellectual property solutions

Altera offers a broad portfolio of easy-to-use IP cores. These IP cores are high-quality “building blocks” that you can drop into your system designs, avoiding the time-consuming task of creating complete designs from scratch. Easy-to-use, pre-verified, and configurable, Altera® IP cores are optimized for the latest Altera devices and are fully supported in Quartus® II design software. Some of the IP cores are offered by leading

third-party IP vendors who have developed, optimized, and qualified their IP products for Altera devices, licensing them directly to our customers.

Discover how quickly you can use these cores to accelerate your system design, lower development costs, and speed your time to market.

Altera IP design flow



Designing with Altera IP

Install	Download and install with Quartus II software.
Parameterize	Use Altera’s powerful, user-friendly MegaWizard® Plug-In tool to parameterize, select the right features for your needs, estimate resources, set up third-party tools, and generate all files necessary for integrating parameterized IP into your design.
Simulate logic in real time	Evaluate all the features of the IP before you buy—including functionality, size, and speed in your system—with the OpenCore® simulation evaluation feature.
Evaluate in hardware	Generate time-limited FPGA programming files using the OpenCore Plus hardware evaluation feature. Then use the files to verify your complete FPGA design on your own hardware, or any other available board—all before you decide to purchase an IP license.
License	When you are completely satisfied with the IP core and are ready to take your design into production, purchase the license and generate the production device programming files. Altera MegaCore® licenses are for perpetual use, support multiple projects, and include upgrades and support for one year. Floating and node-locked licenses are available for all Altera IP. Third-party IP cores are offered under a variety of different licensing terms, conditions, and pricing models. Contact the IP partners directly for details.

Altera IP search engine

Visit our online IP search engine, www.altera.com/ipmegastore (formerly known as the IP MegaStore), for a complete listing of all Altera and third-party partner IP, development kits, and reference designs. The portfolio includes DSP and embedded system cores, interface protocols, and memory controllers as well as a broad range of development kits and application-specific reference designs. Use the search engine to find the right combination of IP and hardware solutions for your next Altera FPGA or ASIC design.

Search engine overview

DSP	IP cores and reference designs used for digital modulation, image, video, and wireless DSP applications. Multiple hardware development kits for faster system integration and verification.
Embedded	Versatile embedded processors with a wide range of peripherals, development kits, and software tools.
Interface protocols	Today's most common system interfaces with development kits for telecom, datacom, and storage applications.
Memory controllers	Extensive portfolio of standard and high-performance memory controllers for use with the suite of FPGA development kits.

Altera megafunction library elements

Another source for Altera IP is the MegaWizard Plug-In within Quartus II design software. You can access an entire library of basic library elements that support more efficient logic synthesis and device implementation. Via parameterization, you can scale the library elements in size and function. As with the individual Altera IP cores listed in the IP search engine, the library elements are optimized for all Altera device architectures. Learn about some of the key megafunction library elements below. Consult the MegaWizard Plug-In for the most complete listing.

Library element	Description
altgx	Designed for use with Altera Stratix® GX and Arria® GX family devices. Enables set-up of one or more transceivers at multiple data rates.
altmempy	Physical interface which comprises silicon features as well as soft logic. It is responsible for the safe transfer of data between the FPGA and the memory, including all aspects of crossing between different clock domains.
Floating-point adder/subtractor [altfp_add_sub]	Implements a floating-point adder/subtractor.
Floating-point multiplier [altfp_mult]	Follows the IEEE 754 standard for floating-point multiplication.
Floating-point divide [altfp_div]	Implements floating-point division functions.
Floating-point square root [altfp_sqrt]	Implements single- and double-precision floating-point square root functions.
Clock control block [altclkctrl]	Implements a basic clock control block.
Memory-based multiplier [altmemmult]	Implements a basic multiplier.
Remote update circuitry [altremote_update]	Enables your designs to take advantage of dedicated remote system upgrade circuitry available in Stratix series devices.
SERDES transmitter/receiver [altlvds]	Implement either an LVDS deserializer receiver or an LVDS serializer transmitter.
PLL [altpll]	Configures the PLLs in the Stratix and Cyclone® series devices.
PLL reconfiguration [altpll_reconfig]	Reconfigures the PLLs in Stratix and Cyclone series devices.
Square root [altsqrt]	Implements the basic square root.
Active serial memory interface [altasmi_parallel]	Implements a basic active serial memory interface (ASMI) block without the need to know the serial interface.
Double data rate (DDR) I/Os [altdio_out, altdio_in, and altdio_bidir]	Implement DDR interface inputs and outputs.

DSP

Product name	Vendor name
Error detection/correction	
Reed-Solomon Compiler, Decoder	Altera Corporation
Reed-Solomon Compiler, Encoder	Altera Corporation
Viterbi Compiler, High-Speed Parallel Decoder	Altera Corporation
Viterbi Compiler, Low-Speed/Hybrid Serial Decoder	Altera Corporation
Filters and transforms	
FFT/inverse FFT (IFFT)	Altera Corporation
CIC Compiler	Altera Corporation
FIR Compiler	Altera Corporation
2D DCT IDCT	Barco Silex
Forward Discrete Cosine Transform (DCT)	CAST, Inc.
Modulation/demodulation	
Numerically Controlled Oscillator Compiler	Altera Corporation
Cable (QAM) Modulator	Commsonic
DVB T Modulator	Commsonic
DVB S2 Modulator	Commsonic
Video and image processing	
Video and Image Processing Suite	Altera Corporation
Color Space Converter	CAST, Inc.
Fast Black and White JPEG Decoder	Barco Silex
Fast Color JPEG Decoder	Barco Silex
JPEG2000 Encoders and Decoders	Barco Silex
JPEG2000 Encoder	Barco Silex
JPEG CODEC	CAST, Inc.
JPEG Encoders and Decoders	CAST, Inc.
Lossless JPEG Encoder and Decoder	CAST, Inc.
Forward Discrete Wavelet Transform (FDWT)	Barco Silex
Inverse Discrete Wavelet Transform (IDWT)	Barco Silex
H.264 Loop Filter	ATEME
H.264 CABAC/CAVLC Bitstream Generator	ATEME
H.264 SD MPEG-4 AVC Encoder	ATEME
MPEG2 Decoder	Amphion Semiconductor, Ltd.

DSP (continued)

Product name	Vendor name
Additional functions	
SHA-1	CAST, Inc.
32- and 128-bit AES Encryption/Decryption Cores	CAST, Inc.
DES Cryptoprocessor	CAST, Inc.
High-Speed AES Encryption/Decryption Cores	D'Crypt Pte. Ltd.
Accelerated Display Graphics Engine	Bitsim
Floating-Point Arithmetic Unit	Digital Core Design
Floating-Point Mathematics Unit	Digital Core Design
Floating-Point Pipelined Divider Unit	Digital Core Design
Floating-Point-to-Integer Pipelined Converter	Digital Core Design
32-/16-bit	
Nios® II Embedded Processor	Altera Corporation
Avalon® Memory-Mapped Components	Altera Corporation
Avalon Streaming Components and Adaptors	Altera Corporation
16-bit Microprocessor, 29116A	CAST, Inc.
ARM® Cortex-M1	Arrow Electronics/ARM
8-/4-bit	
2901 4-bit Microprocessor Slice	CAST, Inc.
8051XC-B Microcontroller	CAST, Inc.
8051XC-A Microcontroller	CAST, Inc.
DR8051 8-bit RISC Microcontroller	Digital Core Design
DR8052EX 8-bit RISC Extended Microcontroller	Digital Core Design
C68000 Microprocessor	CAST, Inc.
CZ80CPU Processor	CAST, Inc.
R80515 Microcontroller	CAST, Inc.
DF6811 CPU 8-bit Microcontroller CPU	Digital Core Design
DFPIC1655X RISC Microcontroller	Digital Core Design

Embedded processors

IP megafunctions (continued)

Interfaces and protocols

Product name	Vendor name
Communication	
8B/10B Encoder/Decoder	Altera Corporation
AAL5	Modelware
Inverse Multiplexing for ATM (IMA) version 1.0/1.1	Modelware
ATM Formatter	Adaptive Micro-Ware, Inc.
ATM Deformatter	Adaptive Micro-Ware, Inc.
CRC Compiler	Altera Corporation
Frame-Mapped GFP Controller	Nuvation
GEOS-10: 10:1 Gigabit Ethernet to SONET Multiplexer	Nuvation
GEOS2+2	Nuvation
Multi-Channel HDLC	Modelware
Single-Channel HDLC	Modelware
POS-PHY Level 2 and 3, Link and PHY	Altera Corporation
POS-PHY Level 4	Altera Corporation
SDLC Controller	CAST, Inc.
SONET/SDH Deframer	Aliathon
SONET/SDH Demapper	Aliathon
SONET/SDH Framer	Aliathon
SONET/SDH Mapper	Aliathon
SPI-4.2 Foundation and Manager	Modelware
T1 Framer	Adaptive Micro-Ware, Inc.
T1 Deframer	Adaptive Micro-Ware, Inc.
UTOPIA Level 2 Master	Altera Corporation
UTOPIA Level 2 Slave	Altera Corporation
Ethernet	
Tri-Speed Ethernet (MAC and PCS)	Altera Corporation
MAC-1G Gigabit Ethernet MAC	CAST, Inc.
Gigabit Ethernet MAC	IFI
High-Performance Gigabit Ethernet MAC	IFI
10/100/1000 1588 Ethernet MAC	MorethanIP
10/100/1000 Ethernet MAC with SGMII	MorethanIP
10/100/1000 Ethernet MAC-Net	MorethanIP
10/100/1000Mbps Full Duplex Ethernet MAC	MorethanIP
AnySpeed Ethernet MAC	MorethanIP
10/100 Ethernet MAC	MorethanIP
10/100 Ethernet MAC	Microtronix
Ethernet Layer 2 Switch	MorethanIP
10 Gigabit Ethernet MAC	MorethanIP
10 Gigabit Ethernet PCS	MorethanIP

Interfaces and protocols (continued)

Product name	Vendor name
High Speed	
Serial RapidIO®	Altera Corporation
Serial RapidIO	Mercury Computer Systems, Inc.
SerialLite II	Altera Corporation
Multi-Gigabit Fibre Channel Transport Core	MorethanIP
HyperTransport™ 8-bit	Altera Corporation
HyperTransport 16-bit	GDA Technologies
1394A Firewire	CAST, Inc.
PCI	
PCI Express Gen1 and Gen2 x1, x4, and x8 Lane	Altera Corporation
PCI Express, Gen1 and Gen2	PLDA
PCI-X Master/Target Core 32-/64-bit	PLDA
PCI Compiler, 32-bit Master/Target	Altera Corporation
PCI Compiler, 32-bit Target	Altera Corporation
PCI Compiler, 64-bit Master/Target	Altera Corporation
PCI Compiler, 64-bit Target	Altera Corporation
32-bit PCI Bus Master/Target Interface	Eureka Technology Inc.
32-bit PCI Bus Target Interface	Eureka Technology Inc.
32-bit PCI Host Bridge	Eureka Technology Inc.
64-bit PCI Bus Master/Target Interface	Eureka Technology Inc.
64-bit PCI Bus Target Interface	Eureka Technology Inc.
64-bit PCI Host Bridge	Eureka Technology Inc.
Integrated PCI Core	Northwest Logic, Inc.
PCI Interface	Northwest Logic, Inc.
32-/64-bit PCI Bus Master/Target Interface, 33/66 MHz	PLDA
32-/64-bit PCI Bus Target Interface, 33/66 MHz	PLDA
32-bit PCI Master/Target Interface	Microtronix/OmniWerks
AMBA-AHB PCI Bridge	PLDA
PCI Bus Arbiter	Eureka Technology Inc.
PCI-ISA Bridge	Eureka Technology Inc.
PCI-PCI Bridge	Eureka Technology Inc.
Serial	
I ² C Bus Controller	CAST, Inc.
I ² C Bus Controller Slave	CAST, Inc.
DI2CM I ² C Bus Interface-Master	Digital Core Design
DI2CSB I ² C Bus Interface-Slave	Digital Core Design
I ² C Master and Slave	Microtronix Inc.
I ² C Master	SLS
I ² C Slave	SLS

IP megafunctions (continued)

Interfaces and protocols (continued)

Product name	Vendor name
Serial (continued)	
C_CAN	Bosch
CAN	CAST, Inc.
Nios_CAN	IFI
Nios II Advanced CAN	IFI
ATA-4 Host Controller	Nuvation
ATA-5 Host Controller	Nuvation
MediaLB Device Interface	IFI
PS2 Interface	SLS
USB High-Speed Function Controller	SLS
USB Full/Low-Speed Function Controller	SLS
CUSB USB Function Controller	CAST, Inc.
CUSB2 USB High-Speed Function Controller	CAST, Inc.
USB 1.1 Host/Device	Microtronix
H16450 and H16450S UARTs	CAST, Inc.
H16550 and H16550S UARTs	CAST, Inc.
D16550 UART with 16 Bytes FIFO	Digital Core Design
H16750 and H16750S UARTs	CAST, Inc.
UART	Eureka Technology Inc.
GPIO8-APB	CAST, Inc.
H8250	CAST, Inc.
Programmable Interval Timer/Counter, 8254	CAST, Inc.
MD5	CAST, Inc.
Smart Card Reader	CAST, Inc.
DSPI Serial Peripheral Interface Master/Slave	Digital Core Design
SD/MMC SPI	El Camino GmbH
SDIO/SD Memory /MMC Host Controller	Eureka Technology Inc.

Interfaces and protocols (continued)

Product name	Vendor name
Serial (continued)	
PowerPC Bus Arbiter	Eureka Technology Inc.
PowerPC Bus Master	Eureka Technology Inc.
PowerPC Bus Slave	Eureka Technology Inc.
PowerPC to PCI Host Bridge	Eureka Technology Inc.
PowerPC/SH/1960 System Controller	Eureka Technology Inc.
Audio and Video	
ASI	Altera Corporation
SD/HD/3G-HD SDI	Altera Corporation
Audio Embed Mux/Demux for SDI	Gennum
I2S Audio CODEC	SLS
DMA	
Scatter Gather DMA Controller	Altera Corporation
DMA Controller	Eureka Technology Inc.
DMA Controller for AHB	Eureka Technology Inc.
Flash	
Compact Flash Interface	SLS
NAND Flash Memory Controller - nflashctrl	CAST, Inc.
ISA/PC Card/PCMCIA/Compact Flash Host Adapter	Eureka Technology Inc.
SDRAM	
DDR and DDR2 SDRAM Controllers	Altera Corporation
DDR and DDR2 SDRAM High-Performance Controllers	Altera Corporation
DDR3 SDRAM High-Performance Controller	Altera Corporation
SDR SDRAM Controller	Northwest Logic, Inc.
SDR/DDR SDRAM Controller	Microtronix
RLDRAM II Controller	Altera Corporation
Pipeline SDRAM Controller	Eureka Technology Inc.
SRAM	
QDR II SRAM Controller	Altera Corporation
QDR II SRAM Controller	Microtronix
FCRAM Controller	Digital Core Design

Memories and memory controllers

FOR MORE INFORMATION

Intellectual property www.altera.com/ipmegastore

Quartus II design software



QUARTUS® II

If you're looking for a design environment that will quickly move you from concept to creation, choose Altera's Quartus® II design software. Number one in performance and productivity for CPLD, FPGA, and ASIC designs, Quartus II software offers complete, automated system definition and implementation, all without requiring lower-level HDL or schematics. This capability—plus its seamless integration with leading EDA software tools and flows—will help turn your ideas into working systems in minutes.

Quartus II software

	Subscription Edition	Web Edition
Devices:	All	MAX® and Cyclone® series; Arria® GX family
Features:	100%	95%
Distribution:	Download and DVD	Download and DVD
Cost:	Paid license	Free license
Operating system support:	Windows, UNIX, Linux	Windows

Quartus II design software features summary

Design flow methodology	Incremental compilation	Improves design timing closure and reduces design compilation times up to 70 percent. Supports team-based design.
	Up-front I/O assignment and validation	Enables PCB layout to begin earlier.
	Pin planner	Eases the process of assigning and managing pin assignments for high-density and high pin-count designs.
	SOPC Builder	Automates adding, parameterizing, and linking IP cores—including embedded processors, coprocessors, peripherals, memories, and user-defined logic.
	Off-the-shelf IP cores	Lets you construct your system-level design using IP cores from Altera's megafunction library and from Altera's third-party IP partners.
	Parallel development of FPGA prototypes and ASICs	Allows for FPGA prototypes and HardCopy® ASICs to be designed in parallel using the same design software and IP.
	Scripting support	Supports command-line operation and Tcl scripting, as well as GUI design processing.
Performance and timing closure methodology	Physical synthesis optimization	Uses post place-and-route delay knowledge of a design to improve performance.
	Design space explorer (DSE)	Increases performance by automatically iterating through combinations of Quartus II software settings to find optimal results.
	Extensive cross-probing	Unmatched support for cross-probing between verification tools and design source files.
	Optimization advisors	Provides design-specific advice to improve design timing performance, resource usage, and power consumption.
	Timing closure floorplan editor	Enables analysis of timing data in the floorplan.
	Chip planner	Reduces verification time (while maintaining timing closure) by enabling small, post place-and-route design changes to be implemented in minutes.
	RTL viewer and technology map viewer	Provides schematic representation that can be used to analyze a design's structure before and after its implementation.
Verification	TimeQuest timing analyzer	Create, manage, and analyze complex timing constraints, and quickly perform advanced timing verification with TimeQuest, an ASIC-strength timing analysis tool with native SDC support.
	SignalTap® II embedded logic analyzer	Supports the most channels, fastest clock speeds, largest sample depths, and most advanced triggering capabilities available in an embedded logic analyzer.
	PowerPlay technology	Enables you to accurately analyze and optimize both dynamic and static power consumption.
	SignalProbe routing	Routes an internal node to an unused or reserved pin for analysis with an external scope or logic analyzer.
Third-party support	EDA partners	Offers EDA software support for synthesis, functional and timing simulation, static timing analysis, board-level simulation, signal integrity analysis, and formal verification.

Free Quartus II Web Edition software

Our free Quartus II Web Edition software for Windows-based PCs includes everything you need to design for Altera's latest CPLD and low-cost FPGA families, as well as our Arria GX FPGA family. Quartus II Web Edition software also includes support for entry-level members of Altera's high-end FPGA families.

Quartus II Web Edition features include:

- Schematic- and text-based design entry
- Integrated VHDL and Verilog HDL synthesis and support for third-party synthesis software
- SOPC Builder system generation software
- Placement, routing, verification, and programming functions
- Mentor Graphics® ModelSim®–Altera® Web Edition simulation software (limited-time offer)
- TimeQuest timing analyzer, an ASIC-strength timing analysis tool
- Installation and evaluation of the Altera MegaCore® IP library
- Support for Windows

Altera subscription program

Altera's subscription program offers a comprehensive suite of premium software and IP products. Included in the subscription are:

- Quartus II Subscription Edition software
 - Support for all Altera devices and software features
 - Support for Windows, UNIX, and Linux operating systems
- Nios® II EDS
- ModelSim–Altera edition simulation software
- Complete IP base suite, which includes full licenses to the following Altera functions:
 - FIR Compiler
 - NCO Compiler
 - FFT Compiler
 - DDR SDRAM Controller
 - DDR SDRAM High-Performance Controller
 - DDR2 SDRAM Controller
 - DDR2 SDRAM High-Performance Controller
 - QDR II SRAM Controller
 - RLDRAM II Controller
 - SerialLite II


DSP Builder



Download the latest software and IP in the Altera subscription program at www.altera.com/download.

DSP Builder is a DSP development tool that interfaces between the Quartus II software and The MathWorks MATLAB/Simulink tools. DSP system design in Altera PLDs requires high-level algorithm and HDL development tools. Altera's DSP Builder integrates these tools by combining the algorithm development, simulation, and verification capabilities of The MathWorks MATLAB and Simulink system-level design tools with VHDL synthesis, simulation, and Altera development tools. DSP Builder shortens DSP design cycles by helping you create the hardware representation of a DSP design in an algorithm-friendly development environment.

DSP design flow overview



DSP Builder features summary

The MathWorks MATLAB and Simulink support	Links The Mathworks MATLAB (signal processing toolbox and filter design toolbox) and Simulink software with Altera Quartus II software.
SignalTap II logic analyzer support	Probes signals from the Altera device on the DSP board and imports the data into the MATLAB work space to facilitate visual analysis.
SOPC Builder and Nios II support	Includes blocks that you can use to build custom logic that works with Nios II embedded processors and other Quartus II SOPC Builder designs.
Testbench generation	Automatically generates a VHDL testbench or Quartus II Vector File (.vec) from MATLAB and Simulink test vectors.
Simulation support	Enables bit- and cycle-accurate design simulation.
Hardware in the loop	Accelerates system-level co-simulation with Simulink and provides advanced debugging features.

SOPC Builder

SOPC Builder is an exclusive Quartus II software tool that enables you to quickly and easily build and evaluate systems at the block level. Using SOPC Builder, you can focus on your custom user logic design, differentiating functions by eliminating manual system integration tasks. In addition to your custom logic, you can select common functions from the Altera or Altera partner IP core libraries to include in your system. SOPC Builder automatically generates interconnect logic to make the system work optimally and creates a testbench to verify functionality, saving valuable design time.

Peripheral expansion of stand-alone processors

SOPC Builder includes a component editor feature so you can easily interface to nearly any external processor or DSP device. If you create an SOPC Builder component interface to your processor, you can add additional I/O pins, prepackaged peripherals, or custom, self-made peripherals in just a few mouse clicks. SOPC Builder will build the system and output header files for your software development team. Your team can then access the peripherals from the external processor using their preferred integrated development environment.

Connection to ASSPs or CPUs via PCI

Many ASSPs and processors include PCI interfaces. Altera's SOPC Builder and the SOPC Builder-Ready PCI Compiler function let you easily build systems that communicate through PCI to these ASSPs or external processors.

SOPC Builder features summary

IP selection and parameter selection	Select and parameterize off-the-shelf IP from Altera and our partners or create your own custom components. Off-the-shelf IP includes the Nios II processor, memory interfaces, common embedded system peripherals, bridges and interfaces, DSP IP, and hardware accelerator peripherals.
System interconnect fabric generation	Uses an optimal interconnect fabric created specifically for the requirements of each system. Integration tasks automatically performed by SOPC Builder include: • Datapath multiplexing between design blocks • Address decoding • Wait-state generation • Dynamic bus sizing • Interrupt priority assignment • Clock domain crossing to connect peripherals or systems operating on different clock domains
Component editor	Allows you to create your own custom SOPC Builder components.
IP reuse	Reuse any custom-created IP core designed for SOPC Builder in future products.
Testbench generation	Outputs testbench suites to test-generated systems.
Header file generation	Outputs a custom header file based on the memory map and components of the generated system.

FOR MORE INFORMATION

Quartus II software
 Quartus II Web Edition and
 Subscription Edition download
 Purchase Quartus II software
 Quartus II literature
 Training

www.altera.com/quartus2
www.altera.com/download
www.altera.com/buysoftware
www.altera.com/literature
www.altera.com/training

Development kits

Altera and our partners offer an array of feature-rich development kits, boards, and companion daughtercards to help you simplify your design process and reduce time to market. These board solutions speed system design by providing you with test and debug platforms for RTL generation, while also allowing you to start developing your application software. With reference designs, cables, and programming hardware, these kits and daughtercards provide an ideal FPGA or CPLD design environment.

Companion daughtercards extend the functionality of Altera's host development kits for a variety of application areas such as video and embedded design. Beginning with our Cyclone® II and Stratix® II series development kits, Altera implemented its newly developed high-speed mezzanine connector (HSMC), enabling greater extensibility and higher performance to these host development kits. The HSMC was conceived by Altera and is manufactured and sold by Samtec. The HSMC specification is available for download from Altera's website.

Our development kits were created to be easy to master, delivering detailed example designs supporting Quartus® II design software and a variety of programming options for Altera® devices. What's more, these kits provide a platform for you to use the Altera IP simulation and hardware evaluation flow. Read on to learn more about the different types of development kits that you can plug into your design process.

DSP development kits

Prototype and debug DSP designs for programmable logic using Altera DSP development kits. These kits include all of the critical design resources you'll need: a DSP development board populated with an Altera FPGA, Quartus II software (one-year, time-limited license), DSP Builder (Quartus II MATLAB/Simulink interface), a 30-day evaluation copy of MATLAB/Simulink, and system reference designs. In addition, by using the hardware evaluation feature of our DSP IP, you can jump-start your designs for video and image, wireless, and other complex digital communications by implementing entire subsystems in hardware within hours.

I/O interconnect development kits

Altera and our partners offer cost-effective development kits to evaluate high-speed interfaces. Consider, as an example, the Transceiver Signal Integrity Development Kit, Stratix II GX Edition. This kit enables you to verify our embedded transceiver for use with a wide variety of interfaces such as PCI Express, Gigabit Ethernet, XAUI, Fibre Channel, Serial RapidIO®, HD-SDI, and other major standards.

Each I/O interconnect board comes with a one-year license to Quartus II design software, a user application, and a library of reference designs. You can also use the development kits and hardware IP evaluation to assess a variety of high-speed interconnect IP cores as well as to rapidly prototype and debug in a real-time environment.

Embedded development kits

For embedded processor system development, Altera's embedded development kits provide a great starting point. The Altera Nios® II development kits include Nios II embedded processors and a perpetual license to create embedded systems, the Nios II EDS, and access to more than 60 peripheral IP cores. The kits contain Quartus II design software (one-year, time-limited license), a feature-rich FPGA-based development board (including power supply, USB-Blaster™ download cable, and LCD display), extensive reference designs, tutorials, and complete documentation. In addition, numerous partners are offering companion boards enabling application-specific embedded Nios II development.

Nios II Development Kit, Stratix Edition



DSP Cyclone III EP3C120 board with daughtercard



General-purpose development kits

Whether designing with our MAX® II CPLDs, Cyclone series FPGAs, or Stratix series FPGAs, you can get off the ground quickly with our portfolio of low-cost, easy-to-use development kits. For example, the MAX II Development Kit provides a complete design environment with all of the software, cables, and accessories needed to evaluate the MAX II feature set or begin prototyping a design before you've received custom hardware. These kits enable increased design productivity whether you're evaluating one of the devices, validating your design, prototyping, or using FPGAs as part of your ASIC prototyping strategy. Tap into the wide variety of general-purpose FPGA and CPLD development kits from Altera and our partners.

Development kits

	Product name and vendor name	Device	Description
DSP	DSP Development Kit, Cyclone III Edition ¹ [DK-DSP-3C120N] Altera Corporation	Cyclone III EP3C120N	This kit is for general DSP or wireless design engineers, regardless of whether you need pre-processing, DSP plus FPGA coprocessing, or post-processing. Complete high-speed analog-to-digital and digital-to-analog conversion capability (16-bit, 200 megasamples/s) is included as well as interfaces to DM642 and DaVinci. Altera's DSP Builder GUI simplifies the information flow between the FPGA toolset and MATLAB/Simulink (30-day evaluation copy included).
	Cyclone III Video and Image Processing Development Kit* Bitec	Cyclone III EP3C120N	This kit is designed to help you start developing complex video applications. It supports various video I/O interfaces, allowing you to get your video data in and out of the Cyclone III FPGA. Different video interfaces are supported using the different daughtercards included in this kit: cards supporting ASI/SDI, composite, component, and digital video interfaces (DVI).
	SPR Development System Bittware	Cyclone III FPGA	The Software Programmable Reconfiguration (SPR) development system provides a system platform to explore software reconfiguration of waveform functionality for high-end signal processing applications such as software-defined radio (SDR). The platform provides a flexible, portable, low-cost environment for SPR development in an AMC and MicroTCA environment, enabling you to quickly and cost-effectively bring your waveform designs to life.
	Video Development Kit, Cyclone II Edition ¹ [DK-VIDEO-2C70N] Altera Corporation	Cyclone II EP2C70N	Composite video input channels (NTSC/PAL), 10-bit BT.656 output, 2 x 2 14-bit, 125 million samples per second (MSPS) A/D, two 14-bit, 165 MSPS D/A, VGA DAC, stereo audio CODEC (96 KHz), EMIF connector, Mictor connector, expansion connector, 256-Mbit DDR2 DIMM, 1-Mbit synchronous SRAM, two EPCS64 devices, MATLAB/Simulink evaluation software, DSP Builder development tool.
	Tetra-PMC Bittware	Cyclone II FPGA	The Tetra-PMC+ (TRPM) board is a high-speed analog input board that provides data capture for four 14-bit A/D channels running at up to 105 MHz. The board streams the data directly to an Altera Cyclone II FPGA, which provides A/D control, data distribution, and front-end processing capabilities. A 32-bit, 66-MHz PCI interface and four TigerSHARC link ports make the captured data available to the host board.
	DSP Development Kit, Stratix II Edition ¹ [DK-DSP-2S60N] Altera Corporation	Stratix II EP2S60N	2-channel 12-bit, 125 MSPS A/D, 2-channel 14-bit 165 MSPS D/A, VGA DAC, stereo audio CODEC 96 KHz, connector for TI C6000 kit, two 40-pin connectors for Analog Devices A/D boards, Mictor connector, RSD232 serial port, RJ45, 32-Mbyte SDR SDRAM, 16-Mbyte flash, 1-Mbit SRAM, 16-Mbyte compact flash, MATLAB/Simulink evaluation software, Quartus II development kit edition (DKE) software, DSP Builder development tool, evaluation DSP IP cores, system reference designs and labs, Nios II reference designs.
	DSP Development Kit, Stratix II Professional Edition ¹ [DK-DSP-2S180N] Altera Corporation	Stratix II EP2S180N	2-channel 12-bit, 125 MSPS A/D, 2-channel 14-bit 165 MSPS D/A, VGA DAC, stereo audio CODEC 96 KHz, connector for TI C6000 kit, two 40-pin connectors for Analog Devices A/D boards, Mictor connector, RSD232 serial port, RJ45, 32-Mbyte SDR SDRAM, 16-Mbyte flash, 1-Mbit SRAM, 16-Mbyte compact flash, MATLAB/Simulink evaluation software, Quartus II DKE, DSP Builder development tool, evaluation DSP IP cores, system reference designs and labs, Nios II reference designs.
	DSP Development Kit, Stratix III Edition [DK-DSP-3SL150N] Altera Corporation	Stratix III EP3SL150	This kit comprises a Stratix III development board with an HSMC equipped with 16-bit analog-to-digital and digital-to-analog converters (operating at up to 200 megasamples/s). The HSMC also has interfaces to TI DSP processors (DM642 and DaVinci), allowing the designs that use Stratix III FPGAs to be created both as stand-alone devices and as companion devices. The kit also contains Altera's Quartus II DKE and DSP Builder software and a 30-day trial of MATLAB/Simulink.

¹ RoHS compliant

Development kits (continued)

	Product name and vendor name	Device	Description
DSP (continued)	Audio Video Development Kit, Stratix II GX Edition ¹ [DK-VIDEO-2SGX90N] Altera Corporation	Stratix II EP2SGX90N	DVI inputs/outputs, SD/HD/SDI inputs/outputs, ASI inputs/outputs, AES3 and S/PDIF audio interfaces, DDR2 DIMM, 2-Mbyte SRAM, HSMC expansion connector, 10/100/1000 Ethernet PHY, 1394 MAC/PHY, USB MAC/PHY, Video and Image Processing Suite evaluation IP cores, SDI reference design.
	B2-AMC Universal Baseband Processing Module BittWare	Stratix II FPGA	Stratix II FPGA, TigerSHARC DSP cluster, full-height, single-wide AMC, configurable to support Serial RapidIO, PCI Express, GbE, and XAUI.
	GT-3U-cPCI CompactPCI Board BittWare	Stratix II GX EP2SGX90	Ruggedized hybrid signal processing, Stratix II FPGA, TigerSHARC DSP cluster, BittWare's ATLANTIS architecture providing 2 Gbps of external I/O throughput, DDR2 SRAM/QDR SDRAM, flash memory.
	GX-AMC Bittware	Stratix II GX FPGA	The GX-AMC (GXAM) is a mid-size, single-wide AdvancedMC that can be attached to Advanced Telecom Compute Architecture (AdvancedTCA) carriers or other cards equipped with AMC bays, and used in MicroTCA systems. The GXAM features a high-density Altera Stratix II GX FPGA, BittWare's ATLANTIS framework (implemented in the FPGA), a front panel I/O interface, a control plane interface via BittWare's FINE interface bridge, an IPMI system management interface, and a configurable 8x SERDES interface supporting a variety of protocols. It also provides 10/100 Ethernet, Gigabit Ethernet, two banks of DDR2 SDRAM, one bank of QDR II SRAM, and flash memory for booting the FPGA and FINE.
	Video Input Daughtercard [DC-VIDEO-TVP5146] Altera Corporation	Daughtercard	Two composite video input channels using TI ADC, support for NTSC/PAL, 10-bit BT.656 output, compatibility with expansion connector found on most Altera development kits; included with Video Development Kit, Cyclone II Edition.
	SC DVI Output Module Bitec	Daughtercard	Supports all Altera development kits with Altera DVI expansion slots.
I/O Interconnect	Arria® GX FPGA Development Kit [EP1AGX60N] Altera Corporation	Arria GX FPGA	Provides a hardware platform for development and testing of high-speed serial interfaces in Arria GX FPGAs. PCIe x4 form factor; includes one HSMC connector, 32-Mbyte DDR2 SDRAM, 512-Mbyte flash, and a PCI Express x4 reference design.
	PCI Development Kit, Cyclone II Edition ¹ [DK-PCI-2C35N] Altera Corporation	Cyclone II EP2C35N	Short-form universal PCI (3.3V or 5.0V) card (32/64 bit, 33/66 MHz), 128-Mbit DDR2 SDRAM, two EPCS64 devices, 100-MHz oscillator, SMA/PCI connector clock input, four user push-button switches, eight DIP switches, eight user LEDs.
	Cyclone II PCI Express Development Kit ¹ [DK-PCI-2C35N] Knott Systems	Cyclone II EP2C35N	Provides a hardware platform for developing and prototyping PCI Express, DDR2 SDRAM, and the 10/100/1000 Ethernet interface.
	1-Lane PCI Express Board HiTech Global	Cyclone II EP2C50	One-lane, low-cost PCI Express development platform.
	4-Lane PCI Express Board HiTech Global	Cyclone II EP2C70	Four-lane PCI Express development board with external PHY.
	PCI Express Design Kit with Cyclone II + External PHY PLDApplications	Cyclone II EP2C20 to EP2C50	Designed to support low-cost Altera FPGA families, includes free board version of PLDA's IP Core for PCIe and a Phillips PCIe x1 PHY.
	PCI Express Development Kit, Stratix II GX Edition ¹ [DK-PCI-2SGX90N] Altera Corporation	Stratix II GX EP2SGX90N	PCI Express IP core (eval), two HSMC connectors, 10/100/1000 Ethernet PHY, two SFP interfaces, 256-Mbit DDR2 SDRAM, 2-Mbit QDR II SRAM, 64-Mbyte flash.
	PCI XpressGXII Board PLDApplications	Stratix II GX EP2SGX90	Includes free PCI Express IP core, x1, x4, x8 configurations, FPGA-based, x8 PCI Express male connector, endpoint PCI Express designs.
	8-Lane PCI Express Board HiTech Global	Stratix II GX EP2SGX90 and EP2SGX130	Eight-lane PCI Express development platform (endpoint).
	Transceiver Signal Integrity Development Kit, Stratix II GX Edition* [DK-SI-2SGX90N] Altera Corporation	Stratix II GX EP2SGX90N	Six full-duplex transceiver channels with SMA connectors, one microstrip channel, four matched stripline channels, 25-MHz and 156.25-MHz clock oscillators, SMA clock input, six user push-button switches, eight DIP switches, eight user LEDs, two 7-segment displays, USB port.

¹ RoHS compliant

Development kits (continued)

	Product name and vendor name	Device	Description
I/O Interconnect (continued)	PCI X/PCI Development Kit (PCIXSYS2) PLDApplications	Stratix II EP2S60 to EP2S180	Ideally suited for ASIC/FPGA prototyping, data acquisition applications.
	DN7000K10PCIS The Dini Group	Stratix II EP2S90, EP2S130, or EP2S180	Complete LE system providing ASIC or IP designers with a vehicle to cost-effectively prototype logic and memory design.
	PCI/PCI-X Board HiTech Global	Stratix II EP2S90, EP2S130, or EP2S180	PCI/PCI-X development platform.
	802.11b Wireless Design Kit Microtronix Inc.	Daughtercard	Kit includes a wireless 802.11b CompactFlash card and a Microtronix CompactFlash board.
	Ethernet USB Expansion Kit Microtronix Inc.	Daughtercard	Kit includes a wireless 802.11b CompactFlash card and a Microtronix CompactFlash board.
	I ² C Design Kit Microtronix Inc.	Daughtercard	Provides an easy way to design, develop, and test the Microtronix I ² C IP core.
	10/100/1000 Ethernet PHY Daughter Board with Marvell PHY MorethanIP	Daughtercard	Provides the ability to implement high-speed Ethernet PHY solutions for prototyping and evaluation and embedded software development.
Embedded	10/100/1000 Ethernet PHY Daughter Board with National Semiconductor PHY MorethanIP	Daughtercard	Provides the ability to implement fast Ethernet solutions for prototyping and evaluation and embedded software development.
	Nios II Embedded Evaluation Kit, Cyclone III Edition ¹ [DK-N2EVAL-3C25N] Altera Corporation	Cyclone III EP3C25N	Beginners can check out the pre-built, eye-catching demos displayed on the LCD touch screen, or do some lightweight development. Advanced microcontroller designers can learn about the “hottest” techniques, multi-processor systems, hardware acceleration using Nios C2H Compiler, or designing a complete system in 30 minutes. The kit includes a complete hardware and software design environment for a 32-bit microcontroller plus FPGA evaluation.
	Cyclone III FPGA Development Kit [PK-DEV-3C120N] Altera Corporation	Cyclone III EP3C120N	8-Mbyte SSRAM, 256-Mbyte DDR2 SDRAM, 64-Mbyte flash, configuration via USB, 10/100/1000 Ethernet and USB ports, on-board oscillators and SMAs, graphics LCD and character LC displays, two HSMC expansion connectors, three HSMC debug cards included, on-board power measurement circuitry. Complete documentation including reference designs: <i>Create your first FPGA design in an hour, Measure Cyclone III FPGA power.</i> This kit also includes Quartus II Web Edition design software, an evaluation edition of Nios processor plus related design suite, and the Altera IP library.
	Nios II Development Kit, Cyclone III Edition ¹ [DK-N2EVAL-3C120N] Altera Corporation	Cyclone III EP3C120N	Altera’s Nios II family of embedded processor-based development kits, the top-selling type of kit sold in the last few years, has been outfitted with the latest in cutting-edge hardware and software technology. The unique combination of high-performance embedded processor power and easy-to-use integrated design software has been updated to take advantage of Cyclone III devices, the industry’s lowest cost, first-to-market 65-nm FPGA family. This development kit provides an ideal environment for developing and prototyping a wide range of price-sensitive, high-performance embedded applications.
	Nios II Development Kit, Cyclone II Edition ¹ [DK-NIOS-2C35N] Altera Corporation	Cyclone II EP2C35N	Perpetual Nios II license, µC-OSII evaluation version, 1-Mbit SRAM, 16-Mbit SDR SDRAM, 16-Mbit flash, CompactFlash connector header, 10/100 Ethernet MAC/PHY, two RS-232, expansion headers, debug Mictor connector, four push-button switches, eight LEDs, two 7-segment LED display.
	Firefly Modules EP2C20, EP2C35, and EP2C50 editions ¹ Microtronix Inc.	Cyclone II EP2C20, EP2C35, and EP2C50	Complete Nios II processor system on a module, up to 95 user I/O available, functionally compatible with Firefly I modules, hardware DSP support.

¹ RoHS compliant

Development kits (continued)

	Product name and vendor name	Device	Description
Embedded (continued)	Desktop NanoBoard NB2DSK01 Altium	Cyclone II FPGA	A LiveDesign-enabled, FPGA-based development board allows rapid and interactive implementation and debugging of FPGA designs and supports multiple target devices in the form of swappable FPGA daughterboards, providing a reconfigurable development platform. Upgraded from the NB1, the Desktop NanoBoard supports Cyclone II devices and provides a much larger number of I/O connections from the target FPGA to the connected peripherals.
	Quicgate Controller/FPGA Dallas Logic	Cyclone II FPGA	Allows implementation of general logic functions and/or Altera Nios II processor operation in a small form-factor module.
	Niomite Controller/FPGA Module Dallas Logic	Cyclone II FPGA	Allows implementation of FPGA logic functions and/or Altera Nios II processor operation in the smallest possible form factor.
	BDK2C35 EBV	Cyclone II FPGA	A broadcast development kit based on Altera's Nios II embedded processor and low-cost Cyclone II FPGA family, this kit includes basic audio/visual (A/V) and IT interface building blocks required to build professional broadcast systems on an FPGA platform while offering extensive customization options via a comprehensive range of add-on cards.
	DBC2C20 EBV	Cyclone II FPGA	An FPGA-based development board designed specifically for Altera's Nios II embedded processors and the .NET framework in the industrial market, this board features a Cyclone II device and is equipped with several industrial I/O standards, such as CAN, RS485, RS232, and 24V I/Os, for direct connection to the industrial automation world. Several IP cores are delivered with the board, including Ethernet MAC, CAN Controller (evaluation), and USB Controller (evaluation).
	Nios II Development Kit, Stratix II Edition ¹ [DK-NIOS-2S60N] Altera Corporation	Stratix II EP2S60N	Perpetual Nios II license, µC-OSII evaluation version, 1-Mbit SRAM, 16-Mbit SDR SDRAM, 16-Mbit flash, CompactFlash connector header, 10/100 Ethernet MAC/PHY RJ45, RS-232, 2 expansion headers, debug Mictor connector, 2 expansion/prototype headers, Quartus DKE, Nios II EDS.
	PARIS automotive development platform TRS-STAR	Stratix II EP2S90 or EP2S180 Optional path to HardCopy® HC210W	Microcontroller development platform to create scalable reference designs for automotive infotainment (head-end, drivers assistance, navigation, etc.); add-on modules and reference designs available. Designs can be ported to Altera's automotive-grade Cyclone III FPGAs or HardCopy ASICs.
	DIGILAB CX II El Camino GmbH	Stratix II EP2S90, EP2S130, or EP2S180	Stratix II development board, stackable with four SSRAM banks, serial flash, security key supported SD/MMC socket, and 56 user I/O pins.
	Lancelot VGA IP Design Kit Microtronix Inc.	Daughtercard	Kit includes a small hardware board with a 24-bit RAMDAC, VGA connector, stereo audio connector, and two PS/2 connectors.
ASIC prototyping	Compact Flash Expansion Kit Microtronix Inc.	Daughtercard	Inexpensive module allows the addition of compact flash cards to the Microtronix Product Starter Kit development board system.
	PROC30M, PROC9M Gidel	Stratix III FPGA	Designed for debug and verification of SOC ASIC designs from 3 to 100+ million gates in size with the ability to run at system clock speeds up to 300 MHz.
	PROCStar II, ProcStar III Gidel	Stratix III, Stratix II FPGA	Provides high-capacity, high-speed, multi-FPGA-based prototyping and end system platforms.
	PROCSuperStar Gidel	Stratix II FPGA	An expandable, building block-type system enabling up to 21 Altera Stratix II EP2S180 FPGAs to be used on the board.
	PROCSpark II Gidel	Cyclone II FPGA	Provides an FPGA-based platform for reconfigurable computing.
	Hpe Midi Gleichmann	Stratix, Stratix II FPGA	A flexible FPGA and ASIC development system to replace the Hpe_compact. The Hpe_midi delivers a low-cost environment, enabling complete system development.
	DN7000K10PCI Stratix-based ASIC Prototyping Kit The Dini Group	Stratix II FPGA	A complete logic emulation system providing ASIC or IP designers with a vehicle to cost-effectively prototype logic and memory design.

* RoHS compliant

Development kits (continued)

	Product name and vendor name	Device	Description
ASIC prototyping (continued)	DNMEG S2GX Stratix II GX-Based ASIC Prototyping Kit The Dini Group	Stratix II GX FPGA	A logic emulation daughtercard enabling ASIC or IP designers to cost-effectively prototype logic and memory designs. The DNMEG_S2GX is hosted on any DN7000 or DN8000 series ASIC Dini Group product, but can also be used alone.
	Cyclone Compute Board Iris Technologies	Cyclone II FPGA	Designed for algorithm acceleration, ASIC prototyping, or simulation acceleration, this board has four Cyclone II FPGAs, a PCIe x4 interface, and 2 Gbytes of DDR2. It can be used either within a system or as a stand-alone unit.
	DSP Compute Board Iris Technologies	Cyclone II FPGA	Designed for signal processing, hardware-in-loop simulation, or coprocessing applications, this board has two Cyclone II FPGAs, two TI 6416T DSPs running at 1 GHz, a PCIe x4 interface, and 1 Gbyte of DDR2 memory. It can be used either within a system or as a stand-alone unit.
	DSP DPA Compute Board Iris Technologies	Cyclone II FPGA	Adds data acquisition to the DSP Compute Board by incorporating two 14-bit 125 megasample/s ADC and two 14-bit 165 megasample/s DACs.
	PROC2S Stratix II FPGA Board GiDEL Limited	Stratix II EP2S60 to EP2S180	Functions as ASIC replacement in customer prototyping environment.
	PROCStar II Development Kit GiDEL Limited	Stratix II EP2S60 to EP2S180 (1–4 devices)	Provides a high-capacity, high-speed FPGA-based platform.
	MAGMA High-Speed Prototyping Board ReFLEX CES	Stratix II EP2S180	High-speed prototyping board featuring four high-density FPGA devices.
	HAC2 Gleichmann Electronics Research	Stratix II EP2S180	A hardware accelerator and cosimulator development tool with key features such as high-speed clock acceleration and hardware-in-the-loop for rapid system exploration.
General purpose	DNMEG S2GX Stratix II GX-Based ASIC Prototyping Kit The DiNi Group	Stratix II GX daughtercard	A logic emulation daughtercard enabling ASIC or IP designers to cost-effectively prototype logic and memory designs.
	Cyclone III FPGA Starter Kit ¹ [DK-START-3C25N] Altera Corporation	Cyclone III EP3C25N	1-Mbyte SSRAM, 16-Mbytes DDR SDRAM, 16-Mbytes parallel flash, configuration via USB, four user push buttons, four user LEDs, power measurement circuitry. Complete documentation including reference designs: <i>Create your first FPGA design in a hour</i> , <i>Measure Cyclone III FPGA power</i> , and <i>Create your first Nios II</i> . This kit also includes Quartus II Web Edition design software, evaluation edition of Nios II processor plus related design suite, and Altera IP library.
	Video Development Kit Bitec	Cyclone III FPGA	This kit contains the Cyclone III EP3C120 Development Kit and two HSMC video interface cards together with a collection of IP cores and reference designs. The kit provides a variety of video interface standards including both digital and analog up to HD resolutions.
	ViClaro III HD Video Enhancement Development Platform Microtronix	Cyclone III FPGA	A video enhancement development platform supporting 100/120-Hz HDTV that is 1080p bandwidth-capable and features 32-bit DDR II SDRAM memory, an HDMI transmitter, an analog/HDMI receiver, and dual LVDS links.
	Cyclone II FPGA Starter Development Kit ¹ [DK-CYCII-2C20N] Altera Corporation	Cyclone II EP2C20N	8-Mbit SDRAM, 512-Kbit SRAM, 1/4-Mbit flash, 24-bit audio CODEC, 10 switches, four push buttons, four 7-segment displays, 10 red and eight green LEDs, VGA, RS-232, and PS/2 ports, two 40-pin expansion ports, SD/MMC socket.
	DB2C5 EBV	Cyclone II EP2C5	Provides an entry point into Altera's current FPGA technology. Three LVDS input channels and output channels are available for high-speed communication with other modules. Available I/O pins are connected with plugs on the underside of the board.
	DB2C20 EBV	Cyclone II EP2C20	FPGA-based development board designed specifically for Altera's Nios II embedded processors and the .NET Framework in the industrial market.
	DE2 Development and Education Board Terasic Technologies Inc.	Cyclone II EP2C35	• Development and education board • Many high-end multimedia features (USB, CD-quality audio, digital television (DTV), Ethernet, SD, and LCD) • Many reference designs with free source code

¹ RoHS compliant

Development kits (continued)

General purpose (continued)

Product name and vendor name	Device	Description
Hpe Mini Gleichmann Electronics Research	Cyclone II EP2C35	An FPGA and ASIC SOC development tool with variants to create simplified or complex CPU systems.
ViClaro II HD Video Enhancement Development Platform Microtronix	Cyclone II FPGA	A platform for quickly creating, testing, and deploying new picture enhancement algorithms. The HDMI transmitter and receiver support the defacto standard HDTV digital interface. The dual LVDS links are 1080p-bandwidth capable and support 50-/60-Hz HDTV display systems. When combined with IP, the board can be used to convert or enhance video quality.
Stratix III FPGA Development Kit [DK-DEV-3SL150N] Altera Corporation	Stratix III EP3SL150	This kit allows rapid and early development of designs that are designed to take advantage of high-performance Stratix III FPGAs. The development board provides general I/Os that connect to on-board switches and indicators, and to the included 2-line LCD and 128x64-pixel graphics display. The board also has non-volatile and volatile memories (64-Mbyte flash, 4-Mbyte pseudo-SRAM, 36-Mbit QDR II SRAM, 128-Mbyte DDR2 DIMM, and 16-Mbyte DDR2 device), HSMC, and triple-speed Ethernet interfaces. The kit is delivered with Quartus II DKE software and all of the cabling required to start using the board straight out of the box.
Hpe Compact Gleichmann Electronics Research	Stratix II EP2S60 to EP2S180	An FPGA and ASIC development tool providing a low-cost environment for creating complete system development.
RUBY II PMC Prototyping Board ReFLEX CES	Stratix II EP2S60 to EP2S180	PCI mezzanine card (PMC) prototyping board supporting Stratix II devices in 1020-pin FPGA packages.
TREX S2 Prototyping System Terasic Technologies Inc.	Stratix II EP2S60 to EP2S180	700 available user I/O pins, flexible and reusable, a selection of various motherboards, high-speed connectors to enable DDR2 memory access.
Stratix II Demo Board Galaxy	Stratix II EP2S180	Designed for IC development and verification, compatible with Altera Stratix II FPGAs in 1020-pin packages.
DB1270 EBV	MAX II EPM1270	Low-cost MAX II CPLD development board providing a hardware platform for a wide range of applications.
MAX II Starter Kit Galaxy	MAX II EPM1270	Kit can also be used by designers interested in simulating and synthesizing programmable logic circuits in their applications.
MAX II Development Kit ¹ [DK-MAXII-1270N] Altera Corporation	MAX II EPM1270N	USB MAC/PHY, PCI Edge connector (3.3V and 5V tolerant), LCD module, SRAM (128K x 8 bit), temperature gauge with serial peripheral interface (SPI), onboard power meter, active I/O sense circuitry, expansion/prototype header, four switches, four LEDs.
DBMAXLED EBV	MAX II EPM1270	Demonstrates that a single MAX II EPM1270 device can drive a large number of PWM channels.
MAX II Micro Terasic Technologies, Inc.	MAX II CPLD	• Development and education board for CPLD design • MAX II micro board with MAX II EPM2210F324C3 (largest CPLD in MAX II series) and on-board USB-Blaster cable • Reference designs with source code
DIGILAB picoMAX Prototyping Board and Starter Kit El Camino GmbH	MAX EPM3032A to EPM7160S	MAX 3000/MAX 7000 starter kit, includes download/programming hardware.
DB3128 EBV	MAX EPM3128A	Low-cost MAX 3000A CPLD development board with 128 macrocells provides an easy entry point into Altera's CPLD technology.
DB3256 EBV	MAX EPM3256A	5.2-megapixel camera daughtercard with selectable frame rates and resolutions.
PM410 StarFabric Compact PCI Carrier Board Parsec	MAX EPM3256A	Two 3.3V PMC sites, 32-/64-bit 33-/66-MHz PCI busses, 2.5-Gbps StarFabric links on J3, supports full PCI bandwidth.

¹ RoHS compliant

Development kits (continued)

	Product name and vendor name	Device	Description
General purpose (continued)	TRDB_DC2 1.3 Megapixel Camera Module Terasic Technologies Inc.	Daughtercard	Complete digital camera reference design with source code in Verilog HDL, user manual with live demo examples; supports exposure, light-control, and motion capture.
	TRDB_LCM Digital Panel Daughtercard Terasic Technologies Inc.	Daughtercard	3.6" digital panel development kit, reference designs (TV player and color pattern generator) with source code in Verilog HDL.
	HSMC DVI Input/Output Module Bitec	Daughtercard	DVI TX/RX module for the HSMC interface enables you to interface their FPGA projects to real-world DVI signals.
	SC DVI Input Module Bitec	Daughtercard	DVI module for the Santa Cruz interface enables you to interface their FPGA projects to real-world DVI signals.
	SC DVI Output Module Bitec	Daughtercard	DVI module for the Santa Cruz interface enables you to drive high-resolution displays with digital clarity.
	SC Camera Bitec	Daughtercard	5.2-megapixel camera daughtercard with selectable frame rates and resolutions.
	SC Proto Bitec	Daughtercard	Prototyping board for the Santa Cruz interface with convenient access points to power and ground with connector break-out.

FOR MORE INFORMATION

Development kits www.altera.com/devkits

Serial configuration devices

Altera® serial configuration devices store the configuration file for Altera's SRAM-based FPGAs, including Arria®, Cyclone®, and Stratix® series FPGAs. We designed our serial configuration devices to minimize cost and board space while providing a dedicated FPGA configuration solution. Because these devices are cost-optimized, they feature fewer pins, smaller packages, and by extension lower prices and performance than parallel configuration solutions.

Key features

- Low cost
- In-system programmability
- Flash memory access
- Small form factor

Serial configuration devices features summary

Low-cost leadership	Serial configuration devices are designed for low cost. They are the ideal complement to the Cyclone FPGA series to create the lowest-cost system-on-a-programmable-chip (SOPC) solution. They also serve as a low-cost, small-form factor solution for the Stratix FPGA series and Arria GX FPGAs.
ISP	ISP increases design flexibility by allowing in-system design updates and reduces costs by streamlining the manufacturing process.
Flash memory access	Several interface peripherals available with the Nios® and Nios II embedded processors allow you to access the serial configuration device as a memory module connected to your embedded system. Device memory capacity not consumed storing configuration data can be used as general-purpose non-volatile memory, which is perfect for program and data storage. You can also use a Nios II processor to modify configuration data, which is useful for in-field system updates.
Small form factor	The serial configuration devices are available in space-saving 8-pin and 16-pin small-outline integrated circuit (SOIC) packages.

Serial configuration devices and memory

Serial configuration device	Memory size (bits)
EPCS1	1,048,576
EPCS4	4,194,304
EPCS16	16,777,216
EPCS64	67,108,864
EPCS128	134,217,728



Serial configuration devices for Cyclone III FPGAs

Cyclone III device	Raw binary file size (bits) ¹	Serial configuration device				
		EPCS1	EPCS4	EPCS16	EPCS64	EPCS128
EP3C5	2,949,120	–	✓	✓	✓	✓
EP3C10	2,949,120	–	✓	✓	✓	✓
EP3C16	4,087,808	–	✓	✓	✓	✓
EP3C25	5,750,784	–	–	✓	✓	✓
EP3C40	9,535,488	–	–	✓	✓	✓
EP3C55	14,893,056	–	–	✓	✓	✓
EP3C80	19,972,096	–	–	✓ ²	✓	✓
EP3C120	28,573,696	–	–	✓ ²	✓	✓

¹ These are uncompressed file sizes.

² This is with the Cyclone III compression feature enabled.

Serial configuration devices for Cyclone II FPGAs

Cyclone II device	Raw binary file size (bits) ¹	Serial configuration device				
		EPCS1	EPCS4	EPCS16	EPCS64	EPCS128
EP2C5	1,265,792	✓ ²	✓	✓	✓	✓
EP2C8	1,983,536	–	✓	✓	✓	✓
EP2C20	3,892,496	–	✓	✓	✓	✓
EP2C35	6,848,608	–	–	✓	✓	✓
EP2C50	9,535,488	–	–	✓	✓	✓
EP2C70	14,319,216	–	–	✓	✓	✓

¹ These are uncompressed file sizes.² This is with the Cyclone II compression feature enabled.

Serial configuration devices for Cyclone FPGAs

Cyclone device	Raw binary file size (bits) ¹	Serial configuration device				
		EPCS1	EPCS4	EPCS16	EPCS64	EPCS128
EP1C3	1,265,792	✓	✓	✓	✓	✓
EP1C4	1,983,536	✓	✓	✓	✓	✓
EP1C6	3,892,496	✓ ²	✓	✓	✓	✓
EP1C12	6,848,608	–	✓	✓	✓	✓
EP1C20	14,319,216	–	✓	✓	✓	✓

¹ These are uncompressed file sizes.² This is with the Cyclone compression feature enabled.

Serial configuration devices for Arria GX FPGAs

Arria GX device	Raw binary file size (bits) ¹	Serial configuration device				
		EPCS1	EPCS4	EPCS16	EPCS64	EPCS128
EP1AGX20C EP1AGX20D	9,640,672	–	–	✓	✓	✓
EP1AGX35C EP1AGX35D	9,640,672	–	–	✓	✓	✓
EP1AGX50C EP1AGX50D	16,951,824	–	–	✓ ²	✓	✓
EP1AGX60C EP1AGX60D EP1AGX60E	16,951,824	–	–	✓ ²	✓	✓
EP1AGX90E	25,699,104	–	–	✓ ²	✓	✓

¹ These are uncompressed file sizes.² This is with the Arria compression feature enabled.

Serial configuration devices for Stratix IV E FPGAs

Stratix IV E device	Raw binary file size (bits) ^{1,3}	Serial configuration device				
		EPCS1	EPCS4	EPCS16	EPCS64	EPCS128
EP4SE100	51,000,000	–	–	–	✓	✓
EP4SE220	102,000,000	–	–	–	✓ ²	✓
EP4SE290	140,000,000	–	–	–	–	✓ ²
EP4SE350	140,000,000	–	–	–	–	✓ ²
EP4SE530	189,000,000	–	–	–	–	✓ ²
EP4SE630	229,000,000	–	–	–	–	✓ ²

¹ These are uncompressed file sizes.

² This is with the Stratix IV compression feature enabled.

³ These values are preliminary.

Serial configuration devices for Stratix IV GX FPGAs

Stratix IV GX device	Raw binary file size (bits) ^{1,3}	Serial configuration device				
		EPCS1	EPCS4	EPCS16	EPCS64	EPCS128
EP4SGX80	51,000,000	–	–	–	✓	✓
EP4SGX100	51,000,000	–	–	–	✓	✓
EP4SGX220	102,000,000	–	–	–	✓ ²	✓
EP4SGX290	140,000,000	–	–	–	–	✓ ²
EP4SGX350	140,000,000	–	–	–	–	✓ ²
EP4SGX530	189,000,000	–	–	–	–	✓ ²

¹ These are uncompressed file sizes.

² This is with the Stratix IV compression feature enabled.

³ These values are preliminary.

Serial configuration devices for Stratix III FPGAs

Stratix III device	Raw binary file size (bits) ^{1,3}	Serial configuration device				
		EPCS1	EPCS4	EPCS16	EPCS64	EPCS128
EP3SL50	22,000,000	–	–	✓ ²	✓	✓
EP3SL70	22,000,000	–	–	✓ ²	✓	✓
EP3SL110	47,000,000	–	–	–	✓	✓
EP3SL150	47,000,000	–	–	–	✓	✓
EP3SL200	66,000,000	–	–	–	✓	✓
EP3SE260	93,000,000	–	–	–	✓ ²	✓
EP3SL340	120,000,000	–	–	–	–	✓
EP3SE50	26,000,000	–	–	–	✓	✓
EP3SE80	48,000,000	–	–	–	✓	✓
EP3SE110	48,000,000	–	–	–	✓	✓

¹ These are uncompressed file sizes.

² This is with the Stratix III compression feature enabled.

³ These values are preliminary.

Serial configuration devices for Stratix II FPGAs

Stratix II device	Raw binary file size (bits) ¹	Serial configuration device			
		EPCS4	EPCS16	EPCS64	EPCS128
EP2S15	4,721,544	✓ ²	✓	✓	✓
EP2S30	9,640,672	–	✓	✓	✓
EP2S60	16,951,824	–	✓ ²	✓	✓
EP2S90	25,699,104	–	✓ ²	✓	✓
EP2S130	37,325,760	–	–	✓	✓
EP2S180	49,814,760	–	–	✓	✓

¹ These are uncompressed file sizes.

² This is with the Stratix II compression feature enabled.

Serial configuration devices for Stratix II GX FPGAs

Stratix II GX device	Raw binary file size (bits) ^{1,3}	Serial configuration device				
		EPCS1	EPCS4	EPCS16	EPCS64	EPCS128
EP2SGX30C EP2SGX30D	9,640,672	–	–	✓	✓	✓
EP2SGX60C EP2SGX60D EP2SGX60E	16,951,824	–	–	✓ ²	✓	✓
EP2SGX90E EP2SGX90F	25,699,104	–	–	✓ ²	✓	✓
EP2SGX130G	37,325,760	–	–	–	✓	✓

¹ These are uncompressed file sizes.

² This is with the Stratix II GX compression feature enabled.

Attend an Altera® technical training course to reduce your time to market and achieve optimal design results. Our courses give you the skills you need to quickly produce high-performance, small-footprint designs. Do you need help deciding which course is right for you? Check out the set of curricula on the web at www.altera.com/training. Each curriculum is designed for a particular audience and provides a flowchart of recommended courses to take.

Altera instructor-led training courses *(All courses are one day in length unless otherwise noted)*

Course category	General description	Course titles
Design languages	Attain the skills needed to design with Verilog HDL and VHDL for programmable logic.	• Introduction to VHDL • Advanced VHDL Design Techniques • Introduction to Verilog HDL • Advanced Verilog HDL Design Techniques
Software	Acquire design entry, compilation, programming, verification, and optimization skills by learning how to use both basic and advanced features of Quartus® II software.	• The Quartus II Software Design Series: Foundation • The Quartus II Software Design Series: Verification • The Quartus II Software Design Series: Timing Analysis • The Quartus II Software Design Series: Optimization
ASICs	Create optimal ASIC designs using an FPGA as a prototyping vehicle. Learn design techniques that take advantage of key device features and resources.	• Designing with HardCopy® II Devices
Embedded systems	Learn to design a Nios® II soft-core microprocessor system in an Altera FPGA and increase software execution speed through hardware acceleration.	• Designing with the Nios II Processor and SOPC Builder (For hardware engineers. Course length: two days.) • Developing Software for the Nios II Processor (For software engineers. Course length: two days.) • Accelerating Software Using the Nios II C2H Compiler
High speed	Implement high-speed I/O protocols using FPGA multi-gigabit transceivers, and implement interfaces to external memory.	• High-Speed Design Using Stratix® II GX Devices • Interfacing to External Memory with Altera FPGAs
DSP	Solve DSP design challenges using Altera technology.	• Implementing DSP Designs in FPGAs

Online offerings

Online classes give you an overview of a variety of features and design techniques. Take advantage of these free, online training modules to preview topics covered in greater depth in our instructor-led classes, brush up on key how-to tips, or even jump-start your design.

Altera free online training courses *(Courses are approximately one hour in length)*

Course category	Course titles
Design languages	• VHDL Basics • Verilog HDL Basics
Software	• Using Quartus II Software: An Introduction • Chinese Version: Using the Quartus II Software: An Introduction • The Quartus II Software Interactive Tutorial • How to Begin a Simple FPGA Design • What's New in the Quartus II Software Version 8.0 • The Quartus II Software Design Series: Foundation (Online Training) (Note: This training is equivalent to the instructor-led course of the same name) • Design Planning Guidelines for High-Density FPGAs • Using Quartus II Software: Schematic Design • Chinese Version: Using the Quartus II Software: Schematic Design • Using Quartus II Software: Simulation • Chinese Version: Using the Quartus II Software: Simulation • Overview of Mentor Graphics ModelSim Software • TimeQuest Timing Analyzer • Chinese Version: Validating Performance with the TimeQuest Static Timing Analyzer • Switching to the TimeQuest Timing Analyzer • Constraining and Analyzing Timing for Source Synchronous Circuits with TimeQuest • Introduction to Incremental Compilation • Best Practices for Incremental Compilation Partitions and Floorplan Assignments - Part 1 of 2 • Best Practices for Incremental Compilation Partitions and Floorplan Assignments - Part 2 of 2 • SignalTap® II Embedded Logic Analyzer • Chinese Version: Design Debugging Using the SignalTap II Logic Analyzer • Using Quartus II Software: Managing Design Changes With Chip Editor • Using Quartus II Software: Chip Planner • Early Pin Planning With Pin Planner in the Quartus II Software • I/O Management • Power Analysis With the Quartus II Software • FPGA to Board Design Flow Using Mentor Graphics Tools • Signal Integrity Analysis With Third-Party Tools
Devices	• Stratix III Devices: Features and Capabilities • Cyclone III Devices: Features and Capabilities • HardCopy II Architecture and Design Flow • HardCopy II Devices: Design Guidelines
Embedded systems	• Developing Software for the Nios II Processor: Tools Overview • Developing Software for the Nios II Processor: Design Flow • Using SOPC Builder • Chinese Version: Using SOPC Builder • Using the Nios II Processor • Chinese Version: Using the Nios II Processor • Japanese Version: System-on-a-Programmable-Chip Design • Using the Nios II Embedded Processor

Online offerings (continued)

Altera free online training courses *(Courses are approximately one hour in length)*


Course category	Course titles
Embedded systems	• Developing Software for the Nios II Processor: Nios II IDE • Developing Software for the Nios II Processor: Debug Primer • Developing Software for the Nios II Processor: HAL Primer • Developing Software for the Nios II Processor: Software Build Flow - Part 1 of 2 • Developing Software for the Nios II Processor: Software Build Flow - Part 2 of 2 • Developing Software for the Nios II Processor: C2H Fundamentals • Japanese Version: Nios II C2H Compiler Fundamentals • Nios II Floating-Point Custom Instructions
Applications	• DSP System Design with DSP Builder • DSP Builder: Migrating from Version 6.1 to 7.1 • Using SOPC Builder • Debugging and Communicating with an FPGA Using the Virtual JTAG Megafunction • Viterbi Decoder • Using Cascaded-Integrator-Comb Filter in Multirate Digital Systems
High speed	• Using High-Performance Memory Interfaces in Altera FPGAs • Transceiver Basics • Implementing PCI Express Solutions in Transceiver-Based FPGAs • Generating a RapidIO MegaCore® Function Variant • Dynamically Reconfiguring Stratix II GX Transceivers • Generating a Triple-Speed Ethernet MegaCore Variant • Serial RapidIO® Design with Stratix IV GX FPGAs • Ethernet Design with Stratix IV GX FPGAs • PCI Express Design with Stratix IV GX FPGAs • Stratix IV GX Transceiver Configuration for Custom Applications
Scripting	• Command Line Scripting • Introduction to Tcl Part 1 of 2 • Introduction to Tcl Part 2 of 2 • Basic Quartus II Software Tcl Scripting Part 1 of 2 • Basic Quartus II Software Tcl Scripting Part 2 of 2

FOR MORE INFORMATION

Training www.altera.com/training

Package dimensions


Legend




For package details, go to www.altera.com/literature/lit-pkg.com

Note: Outermost dimensions are "D" and "E" for both array and peripheral package families.


FineLine BGA (FBGA) and Hybrid FineLine BGA (HFBGA)




Micro FineLine BGA (MBGA)




Ultra FineLine BGA (UFBGA)




Thin Quad Flat Pack (TQFP), Enhanced Thin Quad Flat Pack (EQFP)




Plastic Quad Flat Pack (PQFP), Power Quad Flat Pack (RQFP)




BGA




45.00 x 45.00
1.27




40.00 x 40.00
1.27




35.00 x 35.00
1.27



45.00 x 45.00
1.27



35.00 x 35.00
1.27





35.00 x 35.00
1.27



27.00 x 27.00
1.27

Ordering codes





Product websites and information

- **MAX® CPLD series**

MAX II CPLDs

www.altera.com/max2

MAX 3000A CPLDs

www.altera.com/products/devices/max3k

MAX 7000 CPLDs

www.altera.com/products/devices/max7k

- **Cyclone® low-cost FPGA series**

Cyclone III FPGAs

www.altera.com/cyclone3

Cyclone II FPGAs

www.altera.com/cyclone2

Cyclone FPGAs

www.altera.com/cyclone

- **Arria® GX FPGA family**

Arria GX FPGAs

www.altera.com/arriagx

- **Stratix® high-end FPGA series**

Stratix IV FPGAs

www.altera.com/stratix4

Stratix III FPGAs

www.altera.com/stratix3

Stratix II FPGAs

www.altera.com/stratix2

- **HardCopy® ASIC series**

HardCopy IV ASICs

www.altera.com/hardcopy4

HardCopy III ASICs

www.altera.com/hardcopy3

HardCopy II ASICs

www.altera.com/hardcopy2

- **Nios® II embedded processors**

Nios II processors

www.altera.com/nios2

Nios user forum

www.niosforum.org

Nios II development kits

www.altera.com/nioskits

Nios II C2H Compiler

www.altera.com/c2h

Have comments about this catalog?

Email us at catalog@altera.com.

- **Intellectual property solutions**

Online IP search engine

www.altera.com/ipmegastore

- **Quartus® II design software**

Quartus II software

www.altera.com/quartus2

Quartus II Web Edition and Subscription Edition download

www.altera.com/download

Purchase Quartus II software

www.altera.com/buysoftware

Quartus II literature

www.altera.com/literature

- **Development kits**

Altera® and partner development kits

www.altera.com/devkits

- **Training**

Altera technical training

www.altera.com/training

- **Literature**

More details on Altera and partner offerings

www.altera.com/literature

Wouldn't you like to design without worry?

Wouldn't it be nice if you could focus on your design performance objectives? Without worrying about cost or power dissipation? And without having to balance early adoption of leading-edge technology against risk?

According to Moore's Law, the more transistors we can fit onto a chip, the greater the performance. Following successful 45nm production in 2007, TSMC is the first and only foundry to deliver a 40nm process technology. The goal is to provide advanced technology options that allow designers to differentiate products in today's competitive electronic market places.


The 40nm node features strained silicon, nickel silicide, and copper interconnects. To achieve high-quality yields, we're implementing a combination of 193nm immersion photolithography and extreme low-k (ELK) materials.

Redefining the foundry model

Our longtime partnership with Altera also contributes to our success. From design through tape out, together we're constantly refining the manufacturing process, resulting in chips that are rolled out successfully from node to node.

What's more, we're also redefining the foundry model. Traditionally, foundries and their IC customers worked in their respective areas—manufacturing process development and IC design—somewhat autonomously. This resulted in longer time-to-market and higher costs. Through a new integrated model, TSMC and Altera engineers collaborate on both manufacturing process development and IC design to achieve better and faster results that ultimately benefit you.

TSMC often qualifies its manufacturing process with Altera test chips. Test chips address design issues upfront while validating circuit design and process characteristics. For Altera, a faster manufacturing ramp means a faster production cycle for their devices.



Devices fabricated on 40nm process technologies are robust, reliable, and ready for integration into an array of high-end applications. And, you'll find that they meet your power, performance, and cost expectations.

More for less

TSMC's 40nm process delivers multiple advantages, including:

- A 2.35 times density improvement compared to devices built on 65nm process technology
- A 15% improvement in active power consumption compared to 45nm
- The foundry segment's smallest SRAM cell size at $0.242 \mu\text{m}^2$
- Multiple Vt core devices and 1.2V, 1.8V, and 2.5V I/O options to meet different product requirements

Winning through close collaboration

To fully leverage today's 40nm advanced process, you need a foundry with the proven experience to make you a winner. TSMC's deep, broad collaboration with Altera assures your success from start to finish.



**Altera Corporation**

101 Innovation Drive
San Jose, CA 95134
USA

www.altera.com

Altera European Headquarters

Holmers Farm Way
High Wycombe
Buckinghamshire
HP12 4XF
United Kingdom
Telephone: (44) 1 494 602 000

Altera Japan Ltd.

Shinjuku I Tower 32F
6-5-1, Nishi Shinjuku
Shinjuku-ku, Tokyo 163-1332
Japan
Telephone: (81) 3 3340 9480
www.altera.co.jp

Altera International Ltd.

Unit 11- 18, 9/F
Millennium City 1, Tower 1
388 Kwun Tong Road
Kwun Tong
Kowloon, Hong Kong
Telephone: (852) 2 945 7000
www.altera.com.cn

Copyright © 2008 Altera Corporation. All rights reserved. Altera, the stylized Altera logo, specific device designations, and all other words and logos that are identified as trademarks and/or service marks are, unless noted otherwise, the trademarks and service marks of Altera Corporation in the U.S. and other countries. RapidIO is a trademark of the RapidIO Trade Association. HyperTransport is a trademark of the HyperTransport Consortium. ModelSim is a registered trademark of Mentor Graphics Corporation. All other product or service names are the property of their respective holders. Altera products are protected under numerous U.S. and foreign patents and pending applications, mask work rights, and copyrights. Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services. 35K June 2008

SG-PRDCT-5.0