



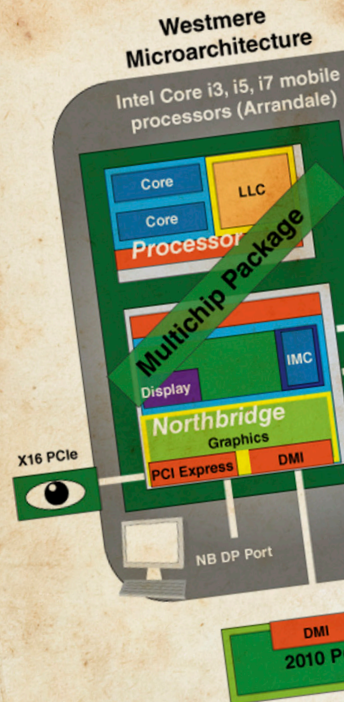



Electronic Design LIBRARY

A compendium of articles from *Electronic Design*

The GRAPHICS CHIP Chronicles Volume 4

Dr. Jon Peddie





CONTENTS



02 Intel i740: Big Investments, Major Embarrassments




CHAPTER 1

03 Microsoft's Talisman: The Graphics Chip That Never Was



CHAPTER 2

07 HP's Artist Graphics Chip: Highly-Integrated With a System Focus




CHAPTER 3

11 The History of the Integrated Graphics Controller



CHAPTER 4

18 Nintendo 64: Breakthrough Design, Genuine Disruption



25 More Resources from Electronic Design



Bill Wong
Editor,
Senior Content
Director

GRAPHICS CONTROLLERS have been an important part of computer systems almost since their inception and they have steadily progressed from providing limited support for low resolution displays to platforms that provide real time ray tracing support. Early display controllers provided features like bit blitting and sprites. Many of these features


started out on PCs to support gaming and other applications. These technologies migrated into cell phones that turned into the smartphones of today that are essentially supercomputers-in-a-pocket.

Staying abreast of the latest and greatest technology, including computer graphics, is what we do at Electronic Design but it is often interesting and useful to look back at what was and how we got to where we are at now. Plus, old technologies and methodologies are often useful now. For example, the 8-bit game programmers moved from PCs to cell phones as the PCs moved onto higher end graphics controllers but the cell phones were more limited at the time. I worked with many of these early graphics chips even before I was the first Lab Director at PC Magazine's famous PC Labs many years ago.

One name that stands out in the graphics world is Jon Peddie who has been following and reporting on this technology area as long as I have. I have known Jon Peddie for many years. He has always been my go-to person when it comes to graphics technology.

Jon has been writing short articles about the graphics chips from a historical perspective and the articles can be found online but we have collected them here so you can easily find them. In this volume we start by examining IBM's XGA.

We hope you enjoy this series of ebooks that delve into the history of graphics.



DR. JON PEDDIE is one of the pioneers of the graphics industry and formed *Jon Peddie Research* (JPR) to provide customer intimate consulting and market forecasting services where he explores the developments in computer graphics technology to advance economic inclusion and improve resource efficiency.

Recently named one of the most influential analysts, Peddie regularly advises investors in the technology sector. He is an advisor to the U.N., several companies in the computer graphics industry, an advisor to the Siggraph Executive Committee, and in 2018 he was accepted as an ACM Distinguished Speaker. Peddie is a senior and lifetime member of IEEE, and a former chair of the IEEE Super Computer Committee, and the former president of The Siggraph Pioneers. In 2015 he was given the Life Time Achievement award from the




CAAD society.

Peddie lectures at numerous conferences and universities worldwide on topics pertaining to graphics technology and the emerging trends in digital media technology, as well as appearing on CNN, TechTV, and Future Talk TV, and is frequently quoted in trade and business publications.

Dr. Peddie has published hundreds of papers, has authored and contributed to no less than thirteen books in his career, his most recent, *Augmented Reality, where we all will live*, and is a contributor to TechWatch, for which he writes a series of weekly articles on AR, VR, AI, GPUs, and computer gaming. He is a regular contributor to IEEE, Computer Graphics World, and several other leading publications.

[👉 BACK TO TABLE OF CONTENTS](#)



CHAPTER 1:

Intel i740: Big Investments, Major Embarrassments

DR. JON PEDDIE

Intel's i740 processor struggled to keep pace with the performance of the competitors such as 3Dfx, Matrox and Nvidia. They rolled out chips that outperformed the i740, forcing Intel to sell the chip at discounts.

Intel has tried several times to get into the stand-alone graphics chip market. Its first attempt in 1982 came with the cross licensing of the NEC 7220, which became the Intel 82720. In 1983, Intel rolled out the iSBX 275 Multibus-based add-in graphics board (AIB) with the chip.

Its second attempt was in 1988 when it released the 82786, which it billed as a VLSI graphics coprocessor. It was designed to be used with Intel's 16-bit 80186 and 80286, and 32-bit 80386 processors. "One of the key hardware extensions that supports the speed needed to do graphics and text is a graphics coprocessor," said Bill Gates at the time. It used VRAM, and Intel said the 82786 can provide virtually unlimited color support and resolution.

In 1989, the company introduced the i860 VLIW RISC processor, code named the N10, which had a 32-bit ALU "Core" along with a 64-bit floating point unit (FPU) that was itself built in three parts: an adder, a multiplier, and a graphics processor. The i860 project was terminated in the mid-1990s and followed by the i960, which was merged with the FPU to become the i960KB and was used in several graphics' terminals.

Intel's i740

One of the more byzantine product developments, however, was Intel's i740. It was started in 1995 when Martin Marietta and Lockheed merged to form Lockheed Martin Corporation, and became the world's largest defense contractor.

Lockheed Martin decided to market the graphics technology they had developed for flight simulators, into civilian use and in January 1995 set up the Real3D division who produced the R3D/100 chip. One of their first customers was Sega. This led to the company's most successful product run, designing the 3D hardware used in over 200,000 Sega Model2 and Model3 arcade systems, two of the most popular systems in history.

In 1997, Intel purchased notebook graphics chipmaker Chips and Technologies for \$430



million. However, no products taking advantage of technology acquired in the merger emerged.

In May 1996, Real3D formed a partnership with Intel and Chips and Technologies to introduce similar technology as an AIB for PCs, a project known as “Auburn.” This project became the AGP-based Intel i740 graphics processor, which Intel released in 1998. Intel also purchased a 20% minority interest in Real3D.

Figure 1. Intel i740 prototype AIB with AGP connector.



The specifications looked great:

Hyper Pipelined Architecture

- 2x AGP Support
- Parallel Data Processing (PDP)
- Precise Pixel Interpolation (PPI)
- Direct Memory Execution (DME)
- Sustained 3D Performance
- 1.1 Meg-Triangles/Second Peak
- 425-500K Triangles/Second Full Featured (66 Mega-Pixel Peak)
- 45-55 Meg-Pixels/Second Full Featured
- Full Sideband Accelerated Graphics Port (AGP) Initiator Support
- Optimized For 440LX Intel AGP sets

3D Graphics Visual Enhancements

- Flat & Gouraud Shading
- Mip-Mapping with Bilinear Filtering (11 LODs)
- Color Alpha Blending for Transparency
- Real Time Texture Paging and Video Texturing
- Fogging & Atmospheric Effects

- Specular Lighting
- Edge Antialiasing
- Stippling Or “Screen Door” Transparency
- Backface Culling
- Z Buffering

3D Graphics Texturing Enhancements

- Per Pixel Perspective Correct Texture Mapping
- Texture Sizes From 1×1 To 1024×1024 pixels (Rect/Sq)
- Texture Formats
- Palletized 1, 2, 4, or 8-bit
- RGBA 1555,565,4444
- Compressed 4:2:2, 555, 1544
- Texture Color Keying
- Texture Chroma Keying
- Integrated Hardware Palette

Local Memory Controller

- 2 to 8 MB SGRAM/SDRAM
- High-Speed 64-Bit 100 MHz Interface

The Intel740 graphics accelerator supported perspective-correct texture mapping, bilinear MIP-Mapping, Gouraud shading, alpha-blending, stippling, antialiasing, fogging and Z Buffering. Its 2D capabilities include BLT and STRBLT engines, a hardware cursor, and an extensive set of 2D registers and instructions.

It also includes dedicated video engines for support of video conferencing and other video applications.

The DME architecture meant that full $2\times$ AGP implementation was integrated into the Intel740 graphics accelerator with sideband operations supporting Type 1, Type 2, and Type 3 sideband cycles. That allowed 533 MB/s peak data transfers. Type 3 support permitted textures to be located anywhere in the 32-bit system memory address space. That was the

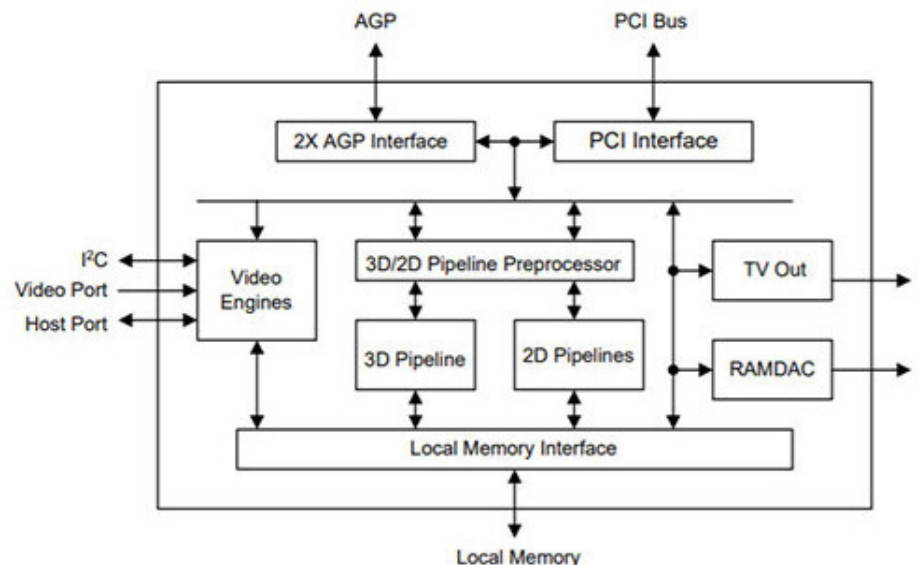


Figure 2. The Intel740 graphics accelerator architectural interfaces.



good news, and the bad, and Intel used the new AGP bus as it had been intended, for texture rendering. Unfortunately texturing from main memory was much slower than from local memory and that approach never caught on—but it would take time to figure that out.

Intel had the brand and the OEM customers from the CPU side of their business. When the company offered the chip to the ODM and OEM suppliers, they went for it. In less than a year, almost a hundred partners had signed up to produce a 740-class AIB and began buying inventory. Intel shipped 4 million 740s between 1997 and 1998, about 4% of the market at the time. Approximately 45 companies showcased 59 boards containing an i740 chip at Computex in 1998.

But the processor couldn't match the performance of the competitors like 3Dfx, ATI, Matrox, or Nvidia. But Intel could more than meet the price, and the other suppliers cried foul. Competitors released chips that outperformed the i740 forcing Intel to sell the chip at low prices to dealers in Asia.

In June 1998, the Federal Trade Commission filed an action against Intel for unfair business practices citing examples of misconduct, and allegedly predatory pricing. The i740 was selling for between \$7 to \$18 in wholesale markets in Taiwan, far below Intel's posted volume price of \$28; the company was accused of dumping. Also, motherboard vendors said it was easier to get Pentium II chips if the vendor also bought i740s.

By late 1999, Intel did two things, shut down the i740 project, and acquired the assets of Real3D from Lockheed Martin. As Real3D crumbled, ATI hired many of the remaining employees and opened up an Orlando office. Prior to the sale of its assets to Nvidia, 3Dfx had sued Real3D over patent infringements. Intel settled the issue by selling all the intellectual property to 3Dfx, which ultimately ended up in Nvidia's hands. Nvidia had SGI's graphics development resources, which included a 10% share in Real3D. That triggered series of lawsuits, joined by ATI. The two companies were involved in lawsuits over Real3D's patents until a 2001 cross-licensing settlement.

Intel exited the discrete graphics chip market for PCs, a market it entered less than 18 months earlier to fanfare and dismal sales. The company continued to produce integrated graphics chipsets, which combine a standard PC chipset with a graphics processor, but those products were targeted at computers selling for \$1,000 and less.

The experience left a bad taste in Intel's mouth and many in the company said the company would never again venture into graphics. Then in 2007 Intel tried again with the Larrabee project. That too ended in failure, and management again said never again. Those people are gone and today in 2019, the company is producing a new generation discrete graphics chip family.


Conclusion

The number 740 turned out to be an attractive designation, and AMD and Nvidia introduced new, unrelated 740 products. In December 2008, AMD announced the RV740 GPU, a 40-nm product. The company started mass production and it proved to be a very successful product.

Then in June 2013 Nvidia introduced the GeForce GT 740M, a DirectX 11 graphics chip targeted at midrange notebooks. It too turned out to be a very popular chip.

to view this article online,  [click here](#)

 **BACK TO TABLE OF CONTENTS**



CHAPTER 2:

HP's Artist Graphics Chip: Highly-Integrated With a System Focus

DR. JON PEDDIE

HP's Artist graphics chip was a highly integrated device designed with the system in mind, combining almost the entire GUI accelerator on a single chip. The only major function not integrated was the frame buffer.

In 1993, facing heavy competition from Sun Microsystems, HP set the design goal for its new 32-bit HP 9000/712 workstation to hit the performance levels of 1992-era workstations and servers at a fraction of their fabrication costs. Their target was the earlier generation HP 9000/735. To accomplish this, HP employed VLSI technologies for the graphics processor components, which were state-of-the-art at the time.

What they came up with was the heart of the system: the famous PA-RISC CPU. The CPU's equally impressive coprocessor was HP's Artist chip, which replaced the company's CRX window accelerator add-in board (AIB). That board marked the beginning of standardized graphics hardware architecture for window system acceleration.

The CRX marked the beginning of a standardized graphics hardware architecture for window system acceleration. The architecture was chosen for its simplicity of implementation and for the clean model it presented to the software driver developers. One of HP's most important design decisions was to accelerate key primitives only—a RISC approach. Many earlier controllers chose to accelerate a large gamut of graphical operations, such as ellipses and arithmetic pixel operations. Graphics subsystems designed with these controllers were typically expensive and exhibited only moderate window system performance. In the CRX and subsequent accelerators, including the Model 712's graphics chip, HP decided to accelerate a smaller but carefully chosen set of primitives, which are described in the following sections.

When the engineers at HP approached the problem of reducing costs, there were three major areas the chip was intended to address (in order of priority):

- Fast 2D GUI
- Digital video decompression support—both locally and over LAN/WAN
- Efficient 3D graphics

In addition, the designers included:

- Vector, rectangle, framebuffer bitBLT, text cursor hardware
- Bit/pixel framebuffer access mode, VRAM block write
- Boolean raster operations
- Two look-up tables to reduce palette conflict

The Artist graphics chip combined a GUI accelerator, a frame buffer controller (32-bits wide), two look up tables (LUTs), video timing, cursor control and an integrated LUT-DAC. The chip was capable of supporting 1 to 2 MB of VRAM and provided 8 bits up to its highest resolution of 1280×1024 at 72 Hz non-interlaced refresh. The ninth bit controlled selection of one of the two LUT-DACs. The chip also included a built-in programmable PLL that eliminated the need for a timing crystal.

The design balanced the CPU's strengths with those of the graphics controller, running the video compression and decompression on the CPU while performing color space conversion and compression/decompression on the Artist graphics chip.




Figure 1. Balanced compression/decompression with CPU and Artist chip. (Image courtesy of HP).

HP also developed a proprietary color compression algorithm that could squeeze 24-bit color to 8 bits while preserving the former's look.

The performance of the Artist chip was impressive at the time:

- Large rectangle fill 850 Mpixels/second
- Vectors/second (10-pixel random) 21 M/second
- 10 x 10 rectangles 1.7M/second
- Text (6 x 13 characters/second) 1M
- 3D transformed vectors/second > 1M
- Frame buffer bitBLT (unaligned pixels/second) 47 M

The CPU handled transformations and clipping as well as lighting, z-buffering and pixel color interpolation for polygons. The Artist chip took care of vector rasterizing and color compression into the frame buffer. The chip added 70 ns VRAMs and included a 37.5 nanosecond (ns) page mode speed, using the features of the VRAM for plane mask, extended data out and block copy. As a result, the chip could deliver 850 Mpixels/second for constant-color objects.

The chip had 525,000 transistors, built in 0.8 micron 3-layer (aluminum) HP CMOS26B process. The die size was 9.7 mm by 12.1 mm and HP enclosed it in 208-pin metal QFP or 240 MQFP packages with a flat-panel driver output. The chip had a 40-80 MHz GUI/RAM

clock and generated a 25-135 MHz video output—and it only used 3.5 watts (W) in the worst case scenario.

Reduced cost was the primary objective for the graphics chip design. And HP succeeded: the manufacturing cost for the Model 712 graphics subsystem was 1/3 of the cost of the original CRX graphics subsystem. In addition, the entry-level 1024-by-768-pixel version of the graphics chip costs five times less than the CRX subsystem.

These cost reductions were achieved primarily through an aggressive amount of integration. The graphics chip represents the culmination of a series of optimizations of the CRX

family, combining almost the entire GUI accelerator onto a single chip. The only major function not integrated was the frame buffer.

The HP Artist chip was one of the first to employ software programmable resolutions. One of the problems with previous workstation graphics subsystems was that they operated at a fixed video resolution and refresh rate. That posed problems in configuring systems at the factory and during customer upgrades.

The Artist graphics chip incorporated an advanced digital frequency synthesizer that generates the clocks necessary for the video subsystem. The synthesizer, which was based on a HP proprietary digital phase-locked loop (PLL) technology, allowed software configurability of the resolution and frequency of the video signal. Thus, alternate monitors could be connected without changing any video hardware. Originally, the chip supported these configurations:

- 640 x 480 pixels 60 Hz, standard VESA timing
- 800 x 600 pixels 60 Hz
- 1024 x 1024 pixels 75 Hz and flat panel
- 1280 x 1024 pixels 72 Hz

As new monitor timings appeared, the graphics chip could be reprogrammed with the parameters associated with the new display.

Summary

HP created its Artist graphics chip with the philosophies of system-level-op-

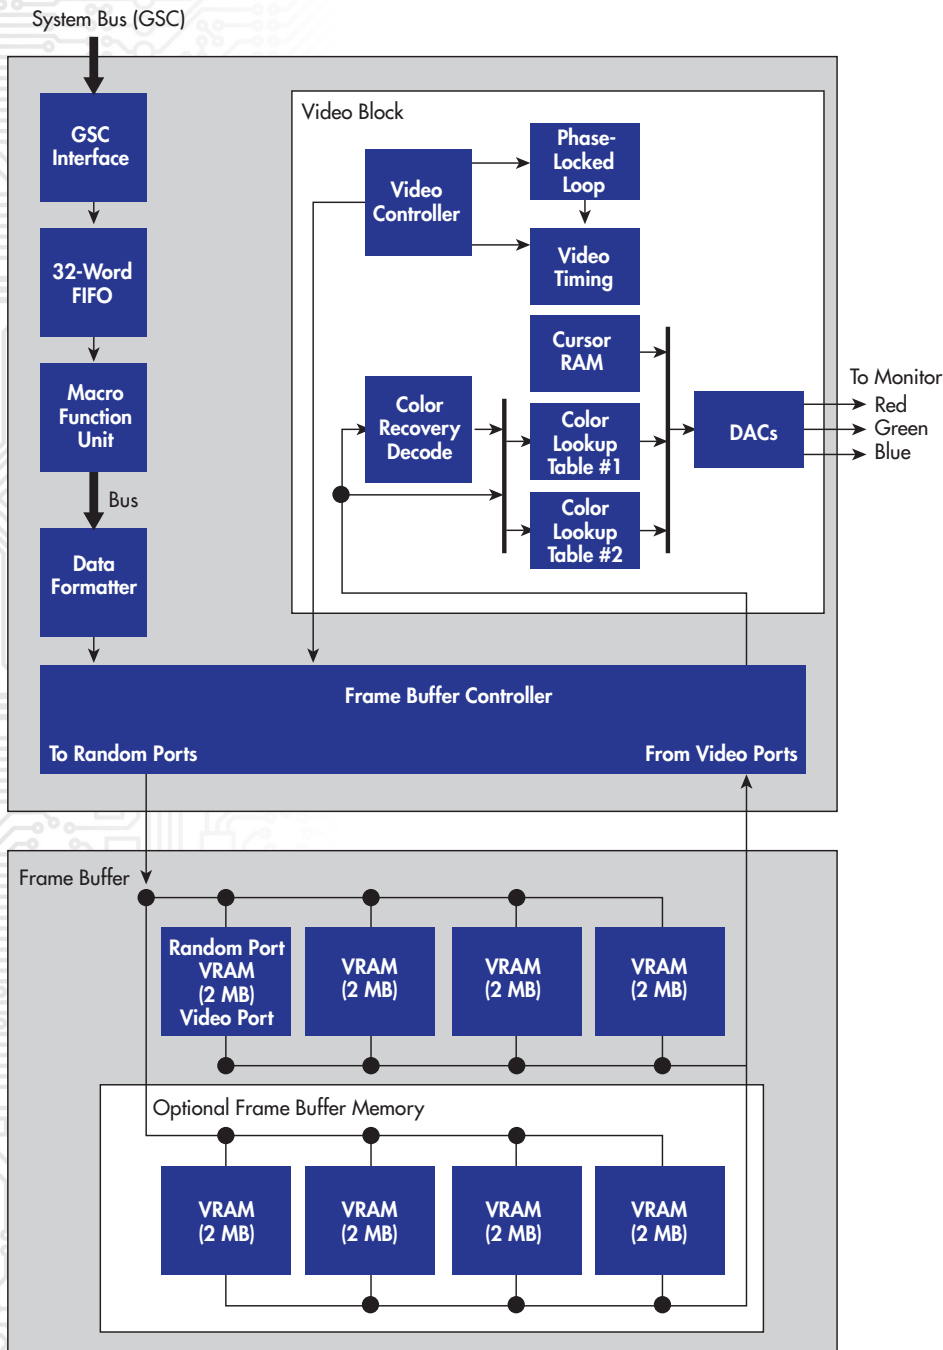



Figure 2. HP's Artist graphics chip block diagram.

timized design. That enabled them to meet their goals of very low manufacturing cost, good performance at their price point, architectural compatibility, and introduction of some important new functionality. The Artist chip was a breakthrough product for HP and served them well for many years.



(Image courtesy of Florida State University's Silicon Zoo).


Conclusion

In the 1980s and early 1990s, semiconductors were laid out on large backlit plotter tables. The layout engineers used to put a signature, or a symbol, a cartoon, or a picture in the pattern somewhere. It was done for three reasons: 1. to sign a work of art; 2. for copyright protection; and 3. because they could. Later, when the layout was done on a computer, many chip designers carried on the tradition.

The above micrograph is a logo that was concealed on an HP chip from the early 1990s. The same chip also featured initials from 20 designers and the hidden message, "If you can read this... you are too damn close!"

to view this article online, [click here](#)

[BACK TO TABLE OF CONTENTS](#)



CHAPTER 3:

The History of the Integrated Graphics Controller

DR. JON PEDDIE

Moore's Law has turned the integrated graphics processing unit (GPU) into one of the key components of personal computers, smartphones, and even automobiles.

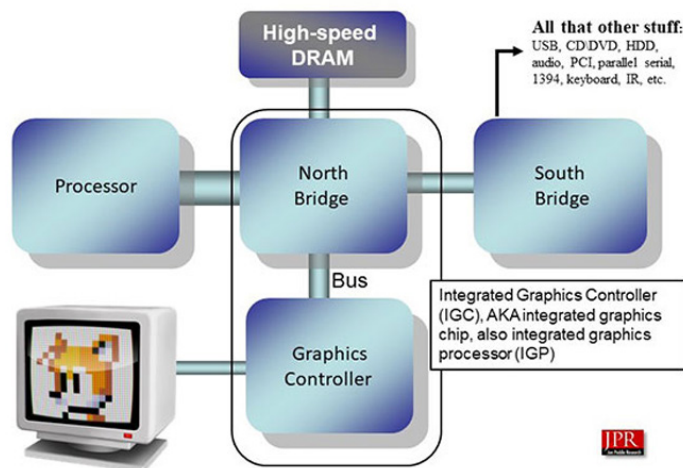
Integrated graphics have evolved from being part of the chipset to being integrated within the CPU. Intel did that first in 2010. AMD followed them with the Llano in 2011, but with a much bigger and more powerful GPU. In between, we saw clever and innovative designs from various suppliers, many of them no longer with us.

May 1995—One of the first examples of integrating a graphics controller with other components was the SPARC enhancement chipset from Weitek. This chipset consisted of two parts: the W8701 SPARC microprocessor and the W8720 Integrated Graphics Controller (IGC). The W8701 integrated a floating-point processor (FPP) into a SPARC RISC microprocessor. It ran at 40 MHz and was socket- and binary-compatible with the SPARC integer unit (IU) standard.

June 1995—Taiwan-based Silicon Integrated Systems introduced the SiS6204, the first PC-based integrated graphics controller (IGC) chipset for Intel processors. It combined the northbridge functions with a graphics controller and set the stage for a new category—the IGC.

SiS developed two IGCs, the 6204 for the 16-bit ISA bus, and the

Basic PC Architecture pre 2010



6205 for the newer PCI bus. The graphics controller offered an integrated VGA with resolution up to $1280 \times 1024 \times 16.8$ million colors (but interlaced), a 64-bit BitBLT engine with an integrated Philips SAA 7110 Video Decoder Interface that provided YUV 4:2:2 support, color-key video overlay support, color space converter, integer video scaling in 1/64th unit increments and VESA DDC1 and DDC2B signaling support. It offered UMA capability with SiS's 551x UMA chipsets. Most importantly, it proved what one could integrate into a small, low-cost chip. SiS and ALi were the only two companies initially awarded licenses to produce third party chipsets for the Pentium 4.

January 1999—In the late 1990s, workstation giant Silicon Graphics Inc (SGI), was trying to meet the threat of the popular and ever-improving x86 processors from Intel. SGI developed the Visual Workstation 320 and 540 workstations using an Intel Pentium processor and designed the Cobalt IGC. It was a massive chip for the time with over 1,000-pins, and it cost more than the standard CPU. It also highlighted the potential performance boost of a unified memory architecture (UMA), one where the graphics processor shared the system memory with the CPU. It made up to 80% of the system RAM available for graphics. However, the allocation was static and only adjusted via a profile.

April 1999—Intel had been leading the industry with consolidating more functions in the CPU. In 1989, when it introduced the venerable 486, it incorporated an FPP, the first chip to do so. A decade later, the company introduced the 82810 IGC (codenamed, Whitney).

September 1999—David Orton, who led the development of the Cobalt chipset while VP of Silicon Graphics' advanced-graphics division, left SGI and became President of ArtX. The company revealed its first integrated graphics chipset with a built-in geometry engine at COMDEX in the fall of 1999, then marketed by Acer Labs of Taiwan. Seeing that, Nintendo contacted ArtX to create the graphics processor (called the Flipper chip) for fourth game console, the GameCube. Then in February 2000, ATI announced it would buy ArtX.

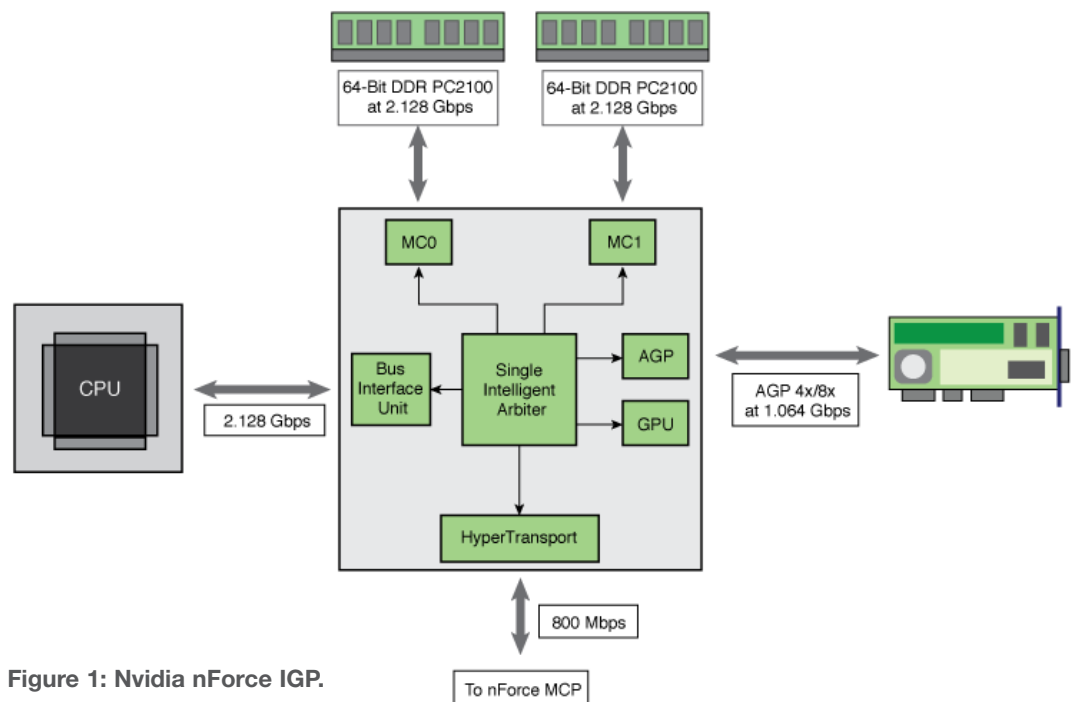


Figure 1: Nvidia nForce IGP.

June 2001—SiS introduced transfer and lighting (T&L) to its IGC. Transformation means producing a two-dimensional view of a three-dimensional scene. Clipping means only drawing the parts of the scene that are in the image after the rendering has completed. Lighting is the process of altering the color of the various surfaces of the scene based on lighting information. Arcade game system boards used hardware T&L since 1993, and it was also used in home video game consoles since the Nintendo 64's Reality Coprocessor GPU (designed and developed by SGI) in 1996. Personal computers implemented T&L in software until 1999.

With the introduction of geometry processing and T&L, the IGC evolved into the IGP — integrated graphics processor.

June 2001—Nvidia introduced its IGP the nForce 220 for AMD Athlon CPU.

The nForce was a motherboard chipset created by Nvidia for AMD Athlon and Duron (later it included support in the 5-series up for Intel processors). The chipset shipped in three varieties; 220, 415, and 420. 220 and 420 were very similar, with each having the integrated GPU.

When Intel moved from a parallel bus architecture to a serial link interface (copying the hyperlink design from AMD), they also declared Nvidia's bus license invalid. After a protracted legal battle, Nvidia won a settlement from Intel, and in 2012 exited the IGP market, leaving only AMD, Intel, and small Taiwanese supplier Via Technologies. Every other company in the market was either bought or driven out of the market by competition.

January 2002—Two years after its acquisition of ArtX, ATI introduced its first IGC, the IGP 320 (code named ATI A3) IGC.

Four years after ATI introduced the IGC, AMD bought ATI to develop a processor with a real integrated GPU. At the time, Dave Orten was ATI's CEO. However, it proved harder

ATI Radeon XPRESS Platform Overview

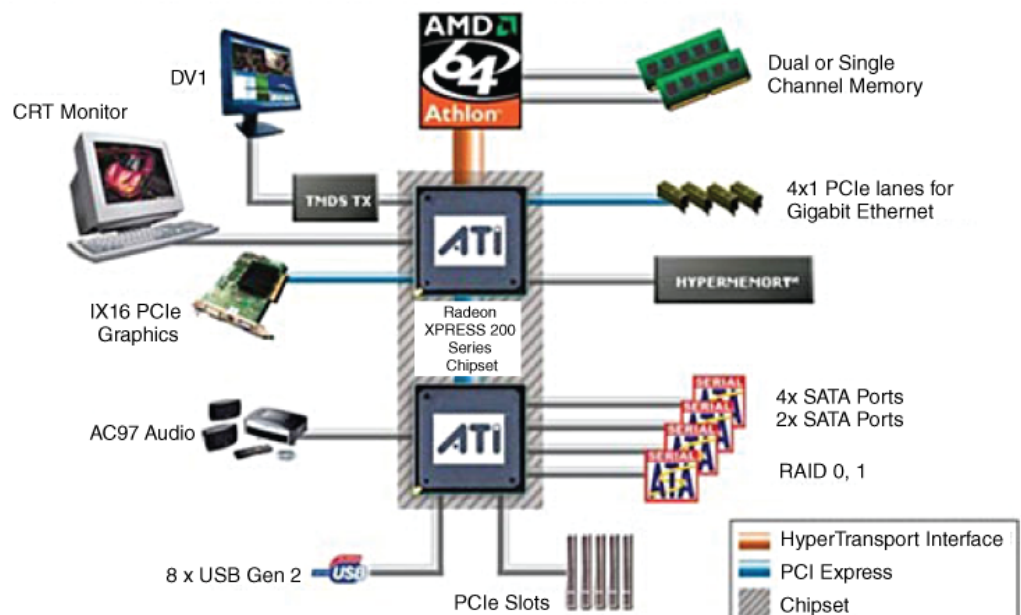


Figure 2: AMD IGP for Athlon processor.

to do than either company thought. Different fabs, different design tools, and clashing corporate cultures made it an incredibly difficult task.

July 2004—Qualcomm introduced its first integrated graphics processor in the MSM6150 and MSM6550 using ATI’s graphics Imageon processor.

The graphics processor could support 100,000 triangles per second and 7 million pixels per second for console-quality gaming and graphics.

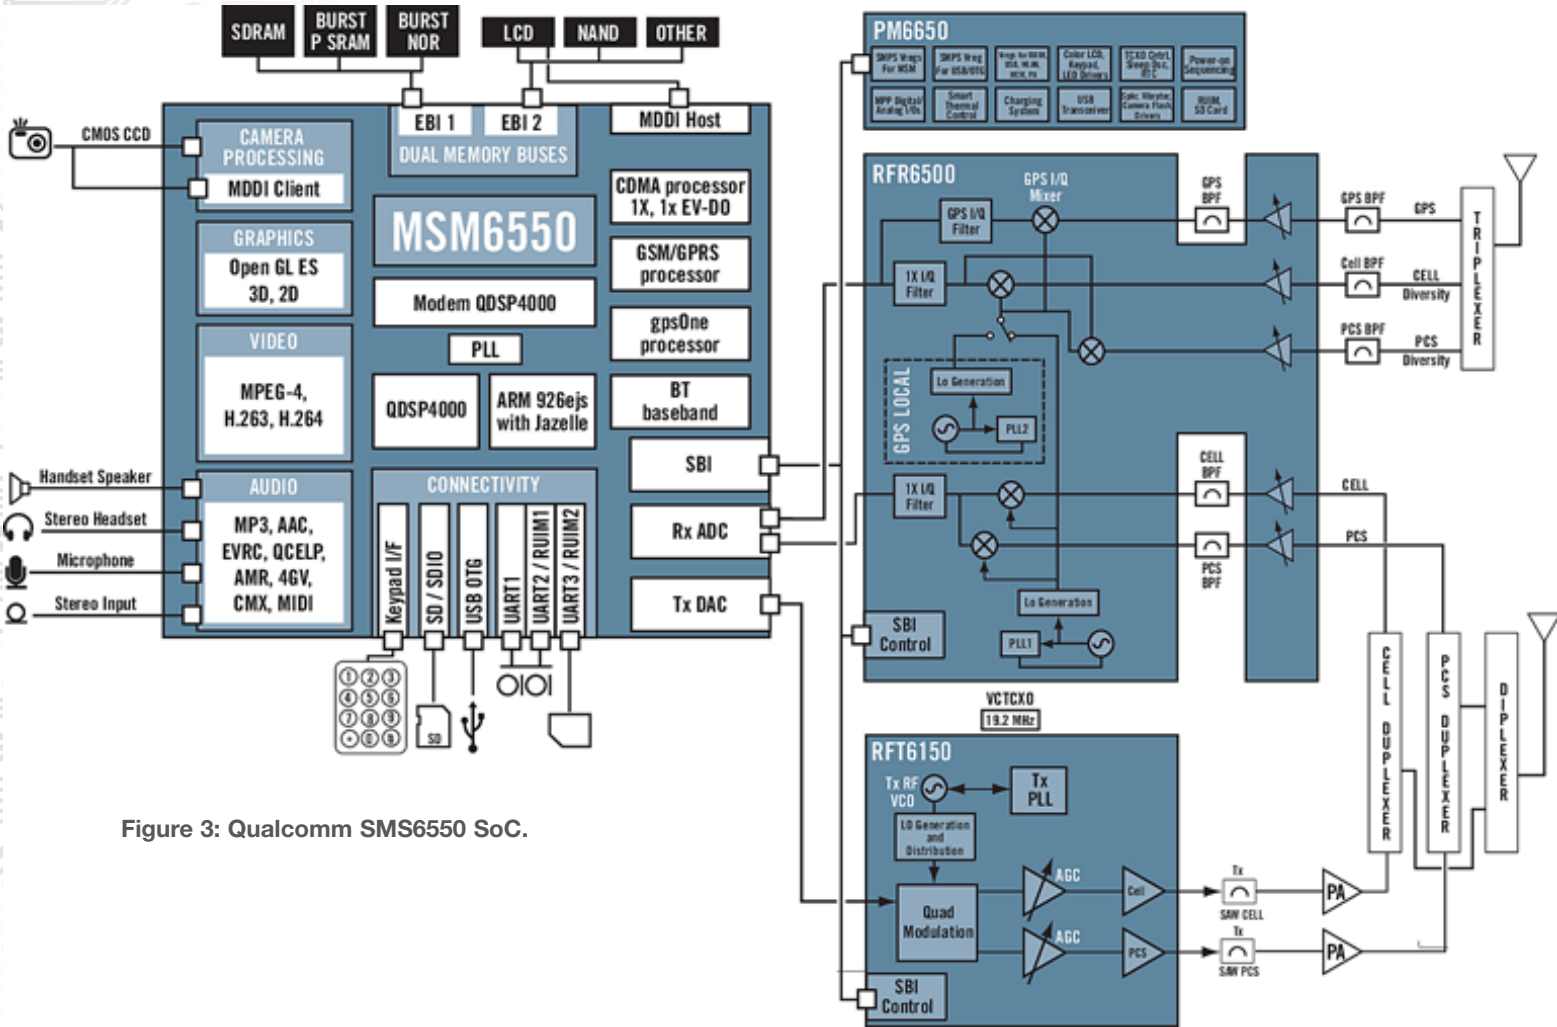


Figure 3: Qualcomm SMS6550 SoC.

October 2005—Texas Instruments introduced the OMAP 2420 and Nokia introduced in the N92 and then the N95.

TI used an Imagination Technologies’ PowerVR GPU design for their OMAP processors. The company was successful with the OMAP in mobile phones until about 2012 when Apple and Qualcomm-based phones took the market by storm.

June and November 2007—Apple introduced the iPhone in the United States in June, and Qualcomm introduced the Snapdragon S1 MSM7227 SoC in November of the same year. Every chip company at that point had developed SoCs with integrated GPUs, primarily for the smartphone market. Apple used Imagination Technologies’ GPU design, and Qualcomm used ATI’s mobile GPU Imageon technology. In January 2009, AMD sold its

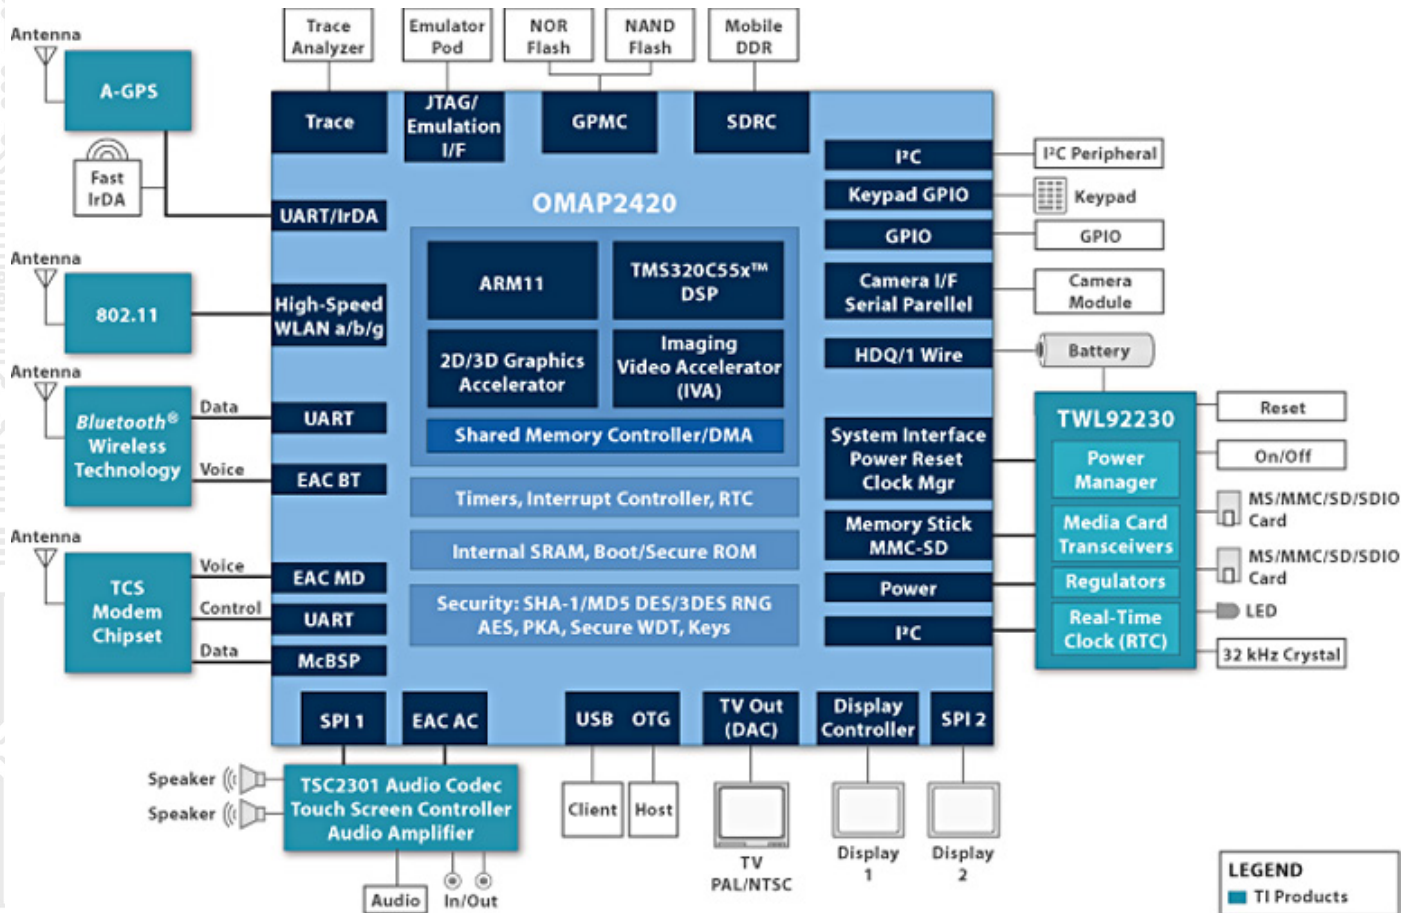


Figure 4: Texas Instrument OMAP 2420 SoC with integrated GPU.

Imageon handheld device graphics division to Qualcomm.

2008—Nvidia introduced the Tegra APX 250 SoC with a 300-to-400 MHz integrated GPU and a 600 MHz ARM 11 processor. Audi incorporated the chip in the entertainment system of its cars, and other car companies followed. In March 2017, Nintendo announced it would use a later generation of the Tegra in its Switch game console.

January 2010—In the PC market, Intel beat out AMD and introduced its Clarkdale and Arrandale processors with Ironlake graphics. Intel branded them as Celeron, Pentium, or Core with HD Graphics. The GPU's specification was 12 execution units (shaders), delivering up to 43.2 GFLOPS running at 900 MHz. The IGP could also decode an H264 1080p video at up to 40 fps.

Intel built the first implementation, Westmere, as a multi-chip product in a single case. The CPU using Intel's 32-nm process and the GPU based on 45-nm.

The most significant difference between Clarkdale and Arrandale is that the latter had integrated graphics. Intel manufactured the fully integrated 131mm² processor, Sandy Bridge, as a 4-core 2.27 GHz processor with the integrated GPU in its 32-nm fab.

January 2011—When AMD bought ATI, Hector Ruiz was president of AMD and Dave Orton was president of ATI. Orton left AMD in 2007, and Ruiz left in 2008, the architects of

Full Integration

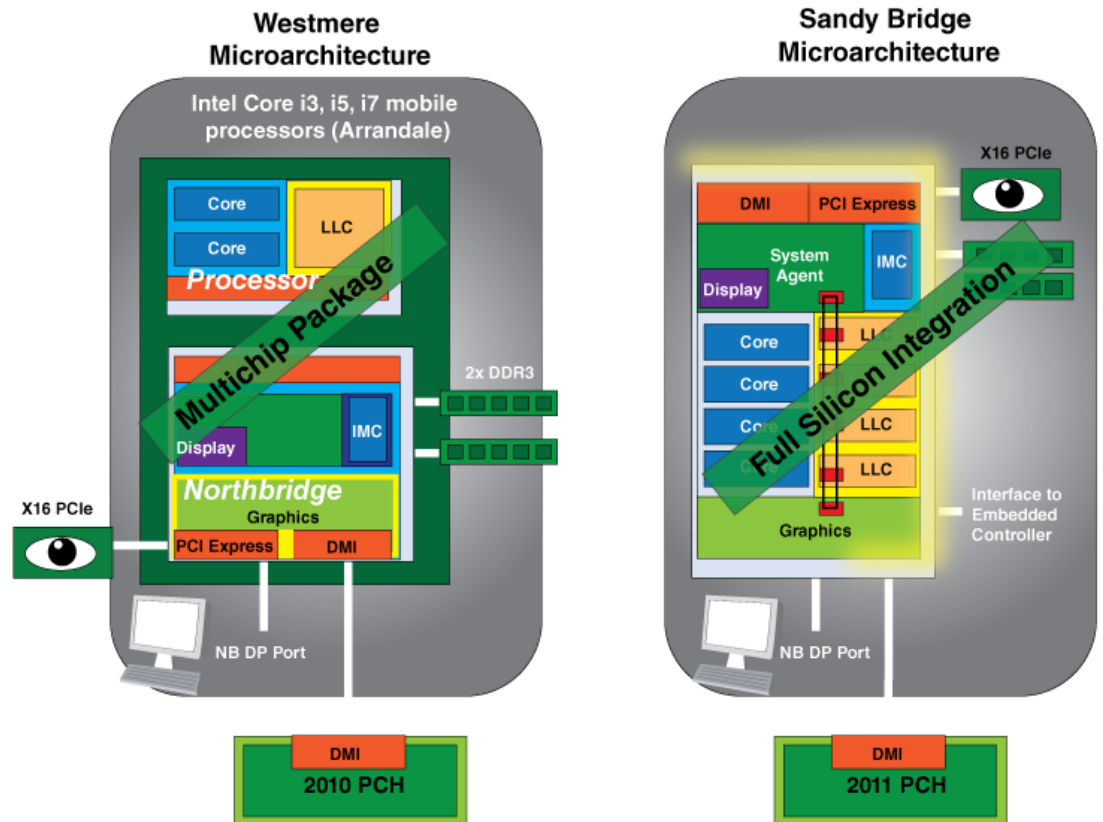


Figure 5: Intel was the first company to incorporate the GPU on the same die as the CPU.

the acquisition and the dream of building a CPU with integrated graphics at AMD. It took three years and several new CEOs after Ruiz left before AMD could introduce an integrated GPU-CPU, which they named an APU—for “accelerated processor unit.” The first product, in 2011, was the Llano, and the internal code name for the processor was Fusion.

The Llano combined the four-core K10 x86 CPU and Radeon HD 6000-series GPU on the same 228mm² die. AMD had it fabricated at Global Foundries with the 32-nm process.

November 2013—Sony introduced PlayStation 4 game console, and Microsoft launched the Xbox One, both based on a custom version of AMD’s Jaguar APU. Sony’s APU used an eight-core AMD x86-64 Jaguar 1.6 GHz CPU (2.13 GHz on PS4 Pro) with an 800 MHz (911 MHz on PS4 Pro) GCN Radeon GPU. Microsoft used an eight-core 1.75 GHz APU (2 quad-core Jaguar modules), and the Xbox One X model contained a 2.3 GHz AMD eight-core APU. The Xbox One GPU ran at 853 MHz, the Xbox One S at 914 MHz, and the Xbox One X at 1.172 GHz using AMD’s Radeon GCN architecture.

Today—The integrated GPU or iGPU is more popular than anything else on the market. It is cost-effective and powerful enough for most graphics chores. It is also seeing acceptance in power-demanding workstation applications.

The iGPU is the dominant GPU used in PCs, and it is in 100% of all game consoles, 100% of all tablets and smartphones, and around 60% of all automobiles, which translates

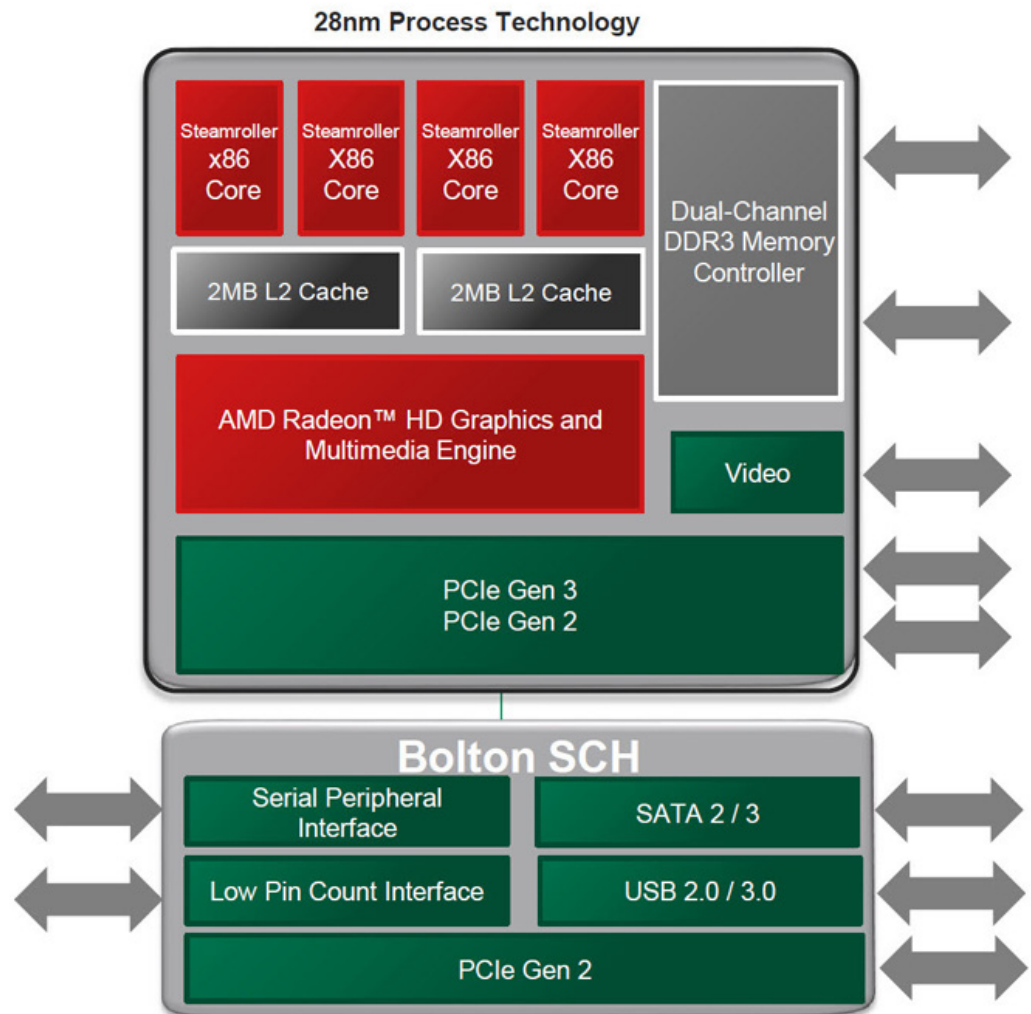



Figure 6: AMD integrated CPU-GPU.

to about 2.1 billion units.

GPUs are incredibly complicated and complex devices with hundreds of 32-bit floating-point processors called shaders and crammed with millions to billions of transistors. That could only have been accomplished due to the miracle of Moore's Law. Every day you engage with multiple GPUs, in your phone, PC, TV, car, watch, your game console, and through the cloud. The world would not have progressed to where it is without the venerable, ubiquitous GPU.

to view this article online, [click here](#)

[BACK TO TABLE OF CONTENTS](#)



CHAPTER 4:

Nintendo 64: Breakthrough Design, Genuine Disruption

DR. JON PEDDIE

In the early 1990s, the idea of shifting state-of-the-art graphics technology from high-end workstations to a consumer product that sold for a couple hundred dollars was seen as a bold and challenging long shot.

Silicon Graphics (SGI) was a leader and highly-respected workstation developer that rose to fame and fortune after it introduced a VLSI geometry processor in 1981. In the following years, SGI developed leading high-end graphics technologies. At that time, an ultra-high-performance workstation could cost more than \$100,000.

Therefore, the idea of adapting such state-of-the-art technology to a consumer product like a video game console that sold for a few hundred dollars was considered bold, challenging, and crazy. Nonetheless, in 1992 and 1993, SGI founder and CEO Jim Clark met with Nintendo CEO Hiroshi Yamauchi to discuss just that—squeezing an SGI graphics system into a console. And thus, the idea of the Nintendo 64 was born.

Even the number, 64, referring to the number of bits, seemed outrageous. Most consoles at the time were struggling to shift from 8-bits to 32 bits. In the video game industry, 64-bits was more or less considered science fiction.

But against all odds, they did it. On November 24, 1995, at Nintendo’s annual Shoshinkai trade show, the




Figure 1: The Nintendo 64 motherboard, CPU, and Reality Coprocessor (RCP), with RMEM under the processors.




Figure 2: Nintendo 64 block diagram.

company revealed the Nintendo 64 console. Then, in May 1996, at the E3 conference in Los Angeles, it showed the Nintendo 64 and announced it would be available in the United States starting in September.

It was an amazing amount of technology crammed into a small package and at an extremely low price of \$250 (roughly \$420 today).

This little gaming supercomputer would be considered feature rich today, and other than the clock speeds, would be a competitive device.

Nintendo 64 features:

- 64-bit custom MIPS R4300 CPU, clock speed of 93.75 MHz
- Rambus DRAM (4 Mbytes) with a maximum bandwidth of 4,500 Mbps
- Sound and graphics, and pixel drawing coprocessors at a clock speed of 62.5 MHz
- Resolutions in the range of 256 × 224 to 640
- The normal resolution is 320 × 240, 24bpp
- 32-bit RGBA frame buffer, with 21-bit color video output
- The graphics processor (RPC) includes:
 - Z-buffer
 - Anti-aliasing
 - Texture mapping: tri-linear interpolated with mip maps, environmental mapping, and perspective correction
- Size: 10.23 inches by 48 inches by 2.87 inches
- Weight: 2.42 pounds
- The system comes with a multifunction 2D and 3D game controller, including digital and analog joysticks, and lots of different buttons.
- MIPS and RPC processors were manufactured by NEC on the 0.35μ process for Nintendo.

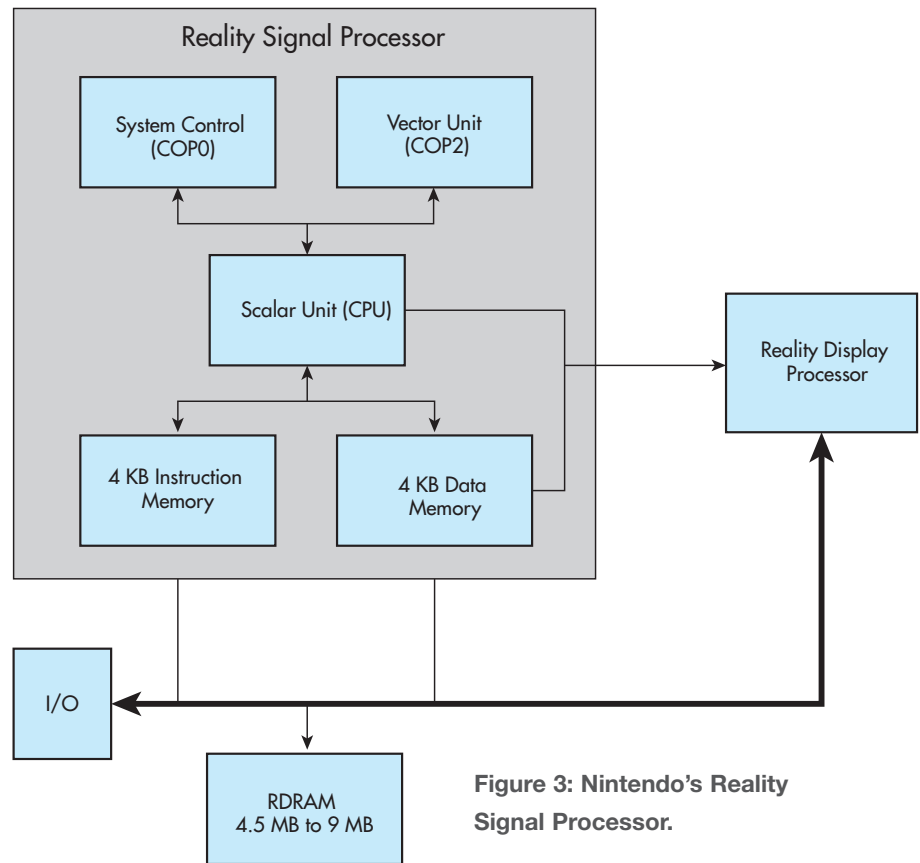


Figure 3: Nintendo's Reality Signal Processor.

The architecture of the system consisted of two major chips: the main CPU and the Reality Coprocessor (RCP) designed by SGI. The following diagram shows the overall arrangement.

The main VR4300 CPU was a 64-bit microprocessor that ran at 93.75 MHz with 64-bit registers, data paths, and buffers to handle high-speed data movement within the chip. The wide data paths were particularly important for operations such as bit-stream decoding and matrix manipulation, core features in video and graphics processing. The VR4300 also supported double-precision floating-point computations for high-performance graphics. Large on-chip caches (16 KB instruction and 8 KB data) delivered high performance for interactive applications by reducing the need for frequent memory accesses. It was built using 0.35 μ NEC process technology.

Within the RCP were two major sub-systems, the Reality Signal Processor (RSP) and the Reality Display Processor (RDP).

Also known as RSP, it contained:

- The Scalar Unit: A MIPS R400-based CPU that implemented a subset of the R400 instruction set.
- The Vector Unit: A co-processor that performed vector operations with thirty-two 128-bit registers. Each register was sliced in eight parts to operate eight 16-bit vectors at once (just like SIMD instructions on conventional GPUs).
- The System Control: Another co-processor that provided DMA functionality and controlled the neighbor display processor module.

To operate the RSP, the CPU stored in RAM a series of commands called a display list along with the data that would be manipulated. The RSP would read the list and apply the required operations on it. The available features included:

- Geometry transformations.
- Clipping and culling (removing unnecessary and unseen polygons).

The RSP fed the Reality Display processor (RDP), which is illustrated in the following block diagram.

After the RSP finished processing polygon data, it sent rasterization commands to the RDP. These commands are either sent using a dedicated bus called XBUS or through main RAM.

The RDP was just another processor with fixed functionality that contained multiple engines used to apply textures over polygons and project them on a 2D bitmap. It could process either triangles or rectangles as primitives. The rectangles are useful for drawing sprites. The RDP's rasterization pipeline contained the following blocks:

- A rasterizer that allocated the initial bitmap that served as a frame-buffer.
- A texture unit that applied textures to polygons using 4 KB of dedicated memory (called TMEM), allowing up to eight tiles to be used for texturing. It could also perform the following operations:

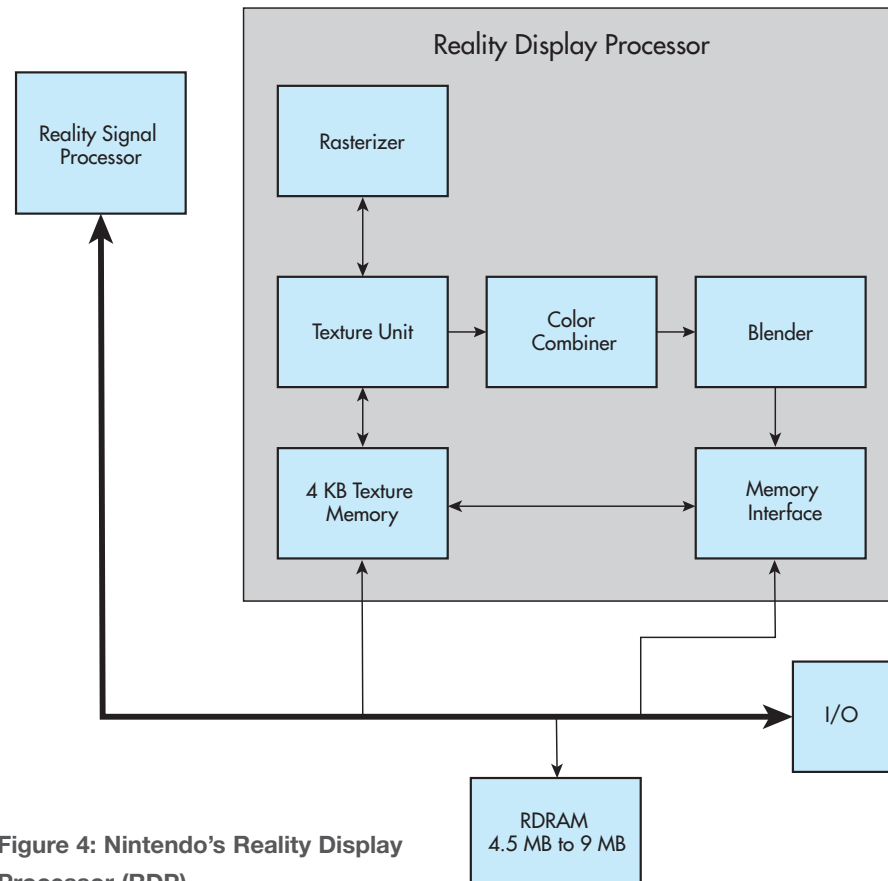


Figure 4: Nintendo's Reality Display Processor (RDP).



- 4-to-1 bilinear filtering for smoothing out textures.
- Perspective correction to improve the coordinate precision of the textures.
- A color combiner that mixed and interpolated multiples layers of colors (for instance, to apply shaders).
- A blender that mixed pixels against the current frame-buffer in order to apply transparency, anti-aliasing, fog, dithering, and z-buffering. That last feature was critical to efficiently cull unseen polygons from the camera viewpoint (replacing software-based polygon sorting methods which could drain a lot of CPU resources).
- A memory interface used by multiple blocks to read and write the current frame-buffer in RAM and/or fill the TMEM.

The RDP provided four modes of functioning, and each mode combined these blocks differently in order to optimize specific operations.

The RDP supported 16.8 million colors. The system could display resolutions from 320 × 240 up to 640 × 480 pixels. Most games tapping into the system's high-resolution 640 × 480 mode required use of the Expansion Pak RAM upgrade.

The system had several advanced, high-end graphics capabilities, including:

- Real-time anti-aliasing—removes jagged edges from the objects, creating a smooth and realistic view as the player moves through a scene.
- Advanced texture mapping techniques—generate high-quality textures and retain the natural texture of every object in the scene, independent of how close the player is to the object.
- Real-time depth buffering—removes hidden surfaces during the real-time rendering process of a scene, allowing game developers to efficiently create 3D environments.
- Automatic load management—enables the objects in the scene to move smoothly and realistically, by automatically tuning the graphics processing.

The console came with a new three-grip controller that allowed 360-degree precision movement. A “3D stick” enabled players to identify any angle in 360 degrees, as well as control the speed of a character's movement. Other new additions include the “C Buttons,” which could be used to change a player's perspective, and a “Z Trigger,” for shooting games.

In addition, the controller featured a memory pack accessory which allowed players to use a special memory card to save game play information on their controller. This enabled players to take their game play data with them and play on other Nintendo 64 systems. Over 350,000 Nintendo 64s sold within days of its release.

Console Wars

The console market was highly contested then, as it is now. New companies were entering the market as older ones were being driven out. As a result, suppliers started a price war that almost ruined them all.

In August 1996, Nintendo announced plans to drop the price of the Nintendo 64 to under \$200 before it launched in the US. The company was looking to match 32-bit systems from Sony and Sega head-to-head on pricing. Sony reacted by reducing prices on many of its video games to \$39.99. Sega, on the other hand, refused to reduce the price of its console.

Then, half a year later, in May 1997, Nintendo announced a new, lower price of \$149.95



for Nintendo 64. At the time, the company cited production efficiencies resulting from a planned uptick in global sales and favorable foreign exchange rates for the price drop. With Sony, Sega, and Nintendo all offering consoles for less than \$150, it was not unreasonable to expect a sub-\$100 sales bonanza for Christmas 1997.

By May 1998, there were only two players in the console market, the Sony PlayStation and the Nintendo 64. The Saturn system was still around—but not for long.

Then, in August 1999, Nintendo dropped the price of the Nintendo 64 to under \$100 as the holiday market heated up.

Conclusion

Lots of advanced techniques and solutions incorporated in the Nintendo 64 have become the basis for modern 3D gaming. The hardware itself had a list of major features:

Here are some of the software features developers played around with that paved the way for modern game engines.

- Trilinear mipmapping, the one most often touted.
- Edge based anti-aliasing, which we have today as FXAA and MLAA.
- Basic real-time lighting, which implies the N64's GPU is a hardware T&L GPU).
- But what stands out in my mind when it comes to the Reality Coprocessor (RCP) is that it was probably one of the first, if not the first, fully programmable GPUs. The processor ran on microcode which developers could tweak to fit their requirements. The problem was Nintendo didn't release developer tools until late in the N64's lifespan. But once they did, several companies (notably Rare and Factor5) pushed the system to its limits.

Other features that were leveraged by game studios included:

- Smart use of clipping. Nintendo favored the use of cartridges for their speed. With clipping, sections of the game world that remained out of the player's sight would not be rendered until the player gets very close to them.
- The game *Banjo Kazooie* introduced a new way to push out large textures for detailed environments. One of the challenges was that it caused memory fragmentation, which means that even though technically enough memory was around, there was not enough contiguous memory to store something in. To solve the problem, the game ran a real-time memory de-fragmenter.
- Texture streaming. This was introduced by developer Factor5 for the game *Indiana Jones and the Infernal Machine*. This technology allowed them to stream textures being rendered, thus overcoming the 4 KB texture memory limit.
- Frame-buffer effects. This is used for effects like motion blur, shadow mapping, "cloaking," and a technique that still amazes me: render to texture. With it, textures are created and updated at runtime.
- Level of detail (LoD). If a model is sufficiently far away, this technique is used to swap it out for a low-poly model.

The Nintendo 64 was truly ahead of its time, but unlike other pioneers the company did not suffer for bringing out such an advanced system. The Nintendo 64 did very well and drove the industry forward toward more realistic and high-performance computer graphics.

Epilogue

Nintendo may have spurred the development of the PlayStation. In the early 1990s,



Nintendo partnered with Sony to develop a new CD-ROM console and attachment for the Super Nintendo system, resulting in a prototype fans called the Nintendo PlayStation.

However, Sony's deal with Nintendo fell through. Sony ultimately decided to ditch Nintendo and launch the PlayStation on its own — a decision that would completely change the course of the video game industry. That ultimately led to the birth of Sony's massive PlayStation brand, and a major and long-term competitor to Nintendo.

In May 1999, Nintendo decided to use IBM's 400-MHz, 128-bit PowerPC chip called Gekko in Nintendo's new Dolphin game console.

Nintendo also said that the new system would use a new graphics chip designed by ArtX, which was founded in 1998 by ex-SGI/MIPS employees who developed the Nintendo 64 graphics processor.

Then, in May 2001, Nintendo announced the GameCube scheduled for launch later the same year. The GameCube featured the Flipper Chip from ATI (through ATI's acquisition of ArtX). The integrated processor incorporated 2D and 3D graphics engines and a DSP for audio processing from Macronix. It also integrated all the system I/O functions, including CPU, system memory, joystick, optical disk, flash card, modem and video interfaces, and an on-chip high bandwidth frame buffer. IBM supplied the 485-MHz Gekko microprocessor.

to view this article online,  [click here](#)

 **BACK TO TABLE OF CONTENTS**


CHECK OUT THESE RESOURCES FROM ELECTRONIC DESIGN AND OUR SISTER BRANDS

WEBSITES

Electronic Design **Microwaves & RF**

sourceesb


Power Electronics



MAGAZINES

You can also apply for a subscription to one of our traditional magazines available in both print and digital formats.


ELECTRONIC DESIGN – complimentary in US and Canada -  Subscribe Now
Non-qualified or Outside the US and Canada: you will be prompted to pay based on your location.

MICROWAVES & RF – complimentary internationally-  Subscribe Now
Non-qualified: you will be prompted to pay based on your location.



NEWSLETTERS

Stay current with the industry with newsletters created by engineers and editors with inside connections! You'll find coverage on a wide variety of electronics technologies from Analog & Mixed Signal to Embedded and more.

 **Click Here** to check out what more than 200,000 engineers are reading now.



ABOUT US

A trusted industry resource for more than 50 years, the Penton Electronics Group is the electronic design engineer's source for design ideas and solutions, new technology information and engineering essentials. Individual brands in the group include *Electronic Design*, *Microwaves & RF* and *Power Electronics*. Also included in the group is a data product for engineers, *SourceToday.com*.

