

SEMI NEWS

AMD launches Carrizo, AMD's sixth-generation APU

Lima to Carrizo in just four years

By Jon Peddie

AMD first started whispering about its forthcoming sixth-generation next-generation, highly integrated mobile A-Series APU, code-named Carrizo, at CES 2015. A bit more was revealed in February and again with AMD's HSA (Heterogenous System Architecture) 1.0 spec roll-out in March. In May at the quaint Clift landmark hotel in San Francisco, AMD finally gave us the details and said shut up until Computex, so we did.

The key features of the new Carrizo are its low power usage, its new CPUs, its graphics, and its video encoder/decoder. It has other features as well, but these are the ones we care about the most.

Graphics

Carrizo has AMD's third generation of Graphics Core Next (GCN, which, by the way, is really kind of a stupid name. How is this graphics core next when they're putting it in the GPU now? Never mind ... marketing), the GCN 1.2, and it is similar to Tonga-based retail AIBs.

AMD increased the GPU's clock speeds and added two more GCN cores in Carrizo from six (384 SPs) to eight (512 SPs), affording a 33% potential improvement, which the company claims gives the APU a 20% improvement in performance, clock for clock.

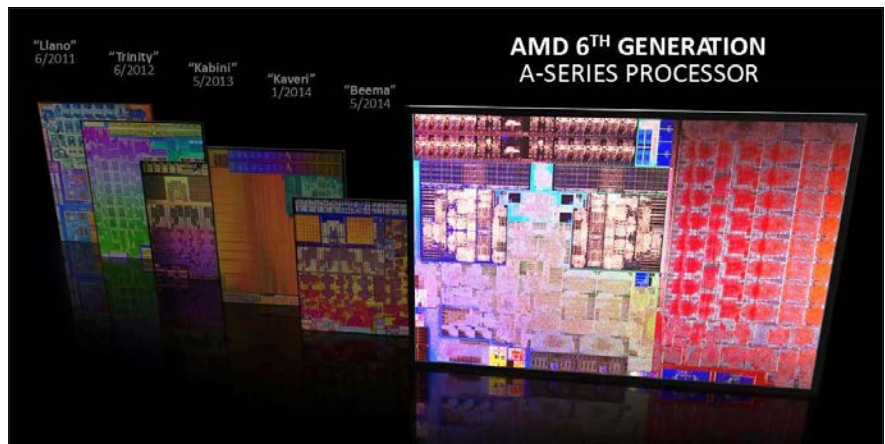
The GCN 1.2 GPU includes AMD's latest color compression algorithms. Part of it is the storing of lossless texture data for maximum visual fidelity and the ease of movement between graphics cores in a compressed state.

GCN 1.2 introduced a new color compression technique for ROPs (render output units) that reduced the memory bandwidth needed for frame buffer operations. AMD has offered color compression for over 10 years since the ATI days; however, GCN 1.2 introduced a color compression method

AMD calls "lossless delta color compression." AMD says it gets a 40% gain in memory bandwidth efficiency with it.

Naturally, the GPU is Mantle- and DX12-compatible. At the time of this writing, we could not find information about support for about Vulkan,

See SEMI NEWS, page 3



SOFTWARE DEVELOPMENTS

Google I/O

Housekeeping with the occasional surprise

By Kathleen Maher

Day 1 of Google I/O was a bit of a chore. The company focused on housekeeping, updates, and going tit for tat with Apple. Among the headline features were Android Pay,



IN THIS ISSUE

- SEMI NEWS: AMD's sixth-gen APU 1
- Broadwell for desktop 6
- Vivante's vision IP 6
- Think Silicon's 3D GPU 7
- SOFTWARE
- Google I/O 1
- Adobe acquires Maximo 9
- Regular features*
- EDITORIAL: Watches 2
- NEWS WATCH 11
- MT. TIBURON TESTING LABS 14
- GUEST EDITORIAL: HDR 22
- TECH INSIDER: Mobile payments 24
- BACK PAGE: The fallacy of TAM 28

new features for watches, and a smarter Google Now—nothing really unexpected, no parachutes, but some nice initiatives that might lead to cool stuff. For instance, Google introduced Jump, a collaboration with GoPro for a 360 camera, and Jacquard, the ability to weave sensors into fabric via conductive fibers.

One of the major themes was the relative openness of Google's interests. The company talked about applications

See SOFTWARE DEVELOPMENTS, page 7

SEMI NEWS

trates the cutting-edge technology available in its GC7000-XS with the addition of vision acceleration to a core that already has photorealistic 3D rendering capabilities.

Vivante says they are working with leading industry partners on high-performance, real-time, intelligent vision image-processing solutions in mobile, automotive, and smart security products through OpenVX- and OpenCL-based Vivante GPUs.

The company claims its OpenVX-compliant GC7000-XS VX allows high-level, deeply embedded devices to optimize their real-time performance for minimal power consumption by allowing continuous sensory vision processes to function without fully powering on the GPU—they say all that in one breath ...

OpenVX can be used directly by applications or to accelerate higher-level middleware, including the popular OpenCV vision library, which is often used for application prototyping. It is also an ideal foundation for use in production systems, with extensive conformance tests to complement a focused and tightly defined specification for consistent and reliable operation across multiple vendors and platforms.

Like other Khronos specifications, OpenVX is extensible to meet individual customer needs before an extension is integrated into the core specification.

■ Think Silicon launches industry's smallest ultra-low power 3D GPU

Tiny GPU IP for IoT

Think Silicon, our favorite GPU IP supplier in Greece, announced their first core product in the NEMA series, the NEMA/t100 design for Internet-of-Things (IoT) graphics applications. The company says they're targeting a new generation of wearable and non-wearable IoT products. The NEMA/t100 is tiny with a silicon footprint of 0.1 mm² (400 MHz in 28 nm). Its low power consumption in leakage and memory power enables longer battery life in wearables. The NEMA/t100 includes OpenGL ES API and Think Silicon's proprietary memory compression technology.

"Now we can truly speak of a tiny GPU where vibrant visual experiences paired with stunning performance can be delivered without the limitations of

Feature	Nema t100	Nema t200	Nema t400
GPU cores	1	2	4
Silicon area (mm ² at 28 nm)	0.1	0.15	0.25
Core clock (MHz at 28 nm)	400	400	400
Shader (GOPs)	3.2	6.4	12.8
Pixel Rate (Mpixels/second)	400	800	1600

TINY 3D GPUS.

reduced battery life or poor power consumption," said Ulli Mueller, vice president of sales and marketing.

The scalable and modular architecture is available in one-, two-, or four-core configurations. The GPU can be customized for small footprint devices and configured for flexible display requirements using core frequencies as low as 25 MHz. Think Silicon features 4 bpp (bits-per-pixel) real-time frame-buffer compression/decompression, and the 6 bpp texture compression and real-time de-compression techniques eliminate the need for external DDR memory.

The NEMA/t100 is available in Verilog HDL code. It supports AMBA (Advanced Microcontroller Bus Architecture) interfaces (AHB, AXI 32 or 64 bits), and embeds DMA controllers with command list for minimal CPU overhead, input/output memory management unit, and an ultra-low-

power network-on-chip. NEMA/t100 supports all major IoT operating systems and middleware including FreeRTOS, Linux, and Android. It comes with software libraries for graphics APIs such as OpenGL ES, DirectFB, and μGFX and a bare metal C library for OS-less systems.

The company's first design win was on a wearable device (smart watch), and the next was on a chipset for white appliances. The idea is that people are used to fluid graphics and nice user interfaces through touch screens. They are happy controlling devices using an interface similar to their smartphone. And using NEMA/t, you can have these features on devices with extremely low silicon/memory/power budget, which you couldn't have until now.

The NEMA/t100 is available for licensing now. ▲

SOFTWARE DEVELOPMENTS

Continued from page 1

for iOS. They announced a Cardboard update that will be able to fit a large 6-inch phone and support iOS content. Likewise, the Unity plug-in for game development will also support iOS platforms. In a similar vein, Nvidia introduced its AndroidWorks set of tools. It's an expansion of Nvidia's GameWorks tool kit, and it supports Android devices beyond those built on the Tegra semiconductor platform. Nvidia's own Shield Android TV set-top box arrived just before Google I/O.

Android M

The housekeeping chores revolved around the Android M developer preview. There was a certain amount of apologizing for Lollipop, we thought. Like every previous release, M promises to improve performance and battery life.

Google is moving their ecosystem forward to the USB-C. Apparently, the

biggest problem with USB is that you had to plug it in right-way-up on earlier versions—and it's impossible to tell which way is the right way on casual examination (see related story, [page 12](#)). USB-C fixes this; you can plug it in either way. Whew, there's a critical technical problem corrected. USB-C is also on the path to becoming a universal connector, and it is able to charge devices much faster than earlier USB versions.

The company promised better battery life with new APIs including sensing to recognize when the device is not being used and dropping down to a deeper sleep to preserve battery life.

Google Now

It's hard not to picture Google Now, Cortana, and Siri all bickering away in a digital space somewhere about who is better. Google has tried to keep Google Now all business with no cute