January 5, 2017

Vega: AMD's New Graphics Architecture for Virtually Unlimited Workloads

Brand new high performance GPU architecture removes many traditional constraints from Gaming, VR, Professional Design and Machine Intelligence arenas

SUNNYVALE, CA -- (Marketwired) -- 01/05/17 -- Today AMD (NASDAQ: AMD) unveiled preliminary details of its forthcoming GPU architecture, Vega. Conceived and executed over 5 years, Vega architecture enables new possibilities in PC gaming, professional design and machine intelligence that traditional GPU architectures have not been able to address effectively. Data-intensive workloads are becoming the new normal, and the parallel nature of the GPU lends itself ideally to tackling them. However, processing these huge new datasets requires fast access to massive amounts of memory. The Vega architecture's revolutionary memory subsystem enables GPUs to address very large data sets spread across a mix of memory types. The high-bandwidth cache controller in Vega-based GPUs can access on-package cache and off-package memories in a flexible, programmable fashion using fine-grained data movement.

"It is incredible to see GPUs being used to solve gigabyte-scale data problems in gaming to exabyte-scale data problems in machine intelligence. We designed the Vega architecture to build on this ability, with the flexibility to address the extraordinary breadth of problems GPUs will be solving not only today but also five years from now. Our high-bandwidth cache is a pivotal disruption that has the potential to impact the whole GPU market," said Raja Koduri, senior vice president and chief architect, Radeon Technologies Group, AMD.

Highlights of the Vega GPU architecture's advancements include:

- ***The world's most advanced GPU memory architecture:*** The Vega architecture enables a new memory hierarchy for GPUs. This radical new approach comes in the form of a new high-bandwidth cache and its controller. The cache features leading-edge HBM2 technology which is capable of transferring terabytes of data every second, doubling the bandwidth-per-pin over the previous generation HBM technology. HBM2 also enables much greater capacity at less than half the footprint of GDDR5 memory. Vega architecture is optimized for streaming very large datasets and can work with a variety of memory types with up to 512TB of virtual address space.
- ***Next-generation geometry pipeline:*** Today's games and professional applications make use of incredibly complex geometry enabled by the extraordinary increase in the resolutions of data acquisition devices. The hundreds of millions of polygons in any given frame have meshes so dense that there are often many polygons being rendered per pixel. Vega's next-generation geometry pipeline enables the programmer to extract incredible efficiency in processing this complex geometry, while also delivering more than 200% of the throughput-per-clock over previous Radeon

architectures.[1] It also features improved load-balancing with an intelligent workload distributor to deliver consistent performance.
- ***Next-generation compute engine:*** At the core of the Vega architecture is a new, next-generation compute engine built on flexible compute units that can natively process 8-bit, 16-bit, 32-bit or 64-bit operations in each clock cycle.[2] These compute units are optimized to attain significantly higher frequencies than previous generations and their support of variable datatypes makes the architecture highly versatile across workloads.
- ***Advanced pixel engine:*** The new Vega pixel engine employs a Draw Stream Binning Rasterizer, designed to improve performance and power efficiency. It allows for "fetch once, shade once" of pixels through the use of a smart on-chip bin cache and early culling of pixels invisible in a final scene. Vega's pixel engine is now a client of the onboard L2 cache, enabling considerable overhead reduction for graphics workloads which perform frequent read-after-write operations.

GPU products based on the Vega architecture are expected to ship in the first half of 2017.

Supporting Resources

- Check out the Vega Preview Cinematic Trailer
- Check out http://ve.ga, the Vega architecture preview portal
- Become a fan of Radeon on Facebook
- Follow @Radeon on Twitter
- Follow @RadeonPro on Twitter
- Follow @RadeonInstinct on Twitter

About AMD
For more than 45 years AMD has driven innovation in high-performance computing, graphics, and visualization technologies -- the building blocks for gaming, immersive platforms, and the datacenter. Hundreds of millions of consumers, leading Fortune 500 businesses, and cutting-edge scientific research facilities around the world rely on AMD technology daily to improve how they live, work, and play. AMD employees around the world are focused on building great products that push the boundaries of what is possible. For more information about how AMD is enabling today and inspiring tomorrow, visit the AMD (NASDAQ: AMD) website, blog, Facebook and Twitter pages.

AMD, the AMD Arrow logo, Radeon and combinations thereof, are trademarks of Advanced Micro Devices, Inc. Other names are for informational purposes only and may be trademarks of their respective owners.

Cautionary Statement

This press release contains forward-looking statements concerning Advanced Micro Devices, Inc. (AMD) including the features, functionality and expectations regarding AMD's Vega GPU architecture and the timing and availability of GPU products based on AMD's Vega architecture, which are made pursuant to the Safe Harbor provisions of the Private Securities Litigation Reform Act of 1995. Forward-looking statements are commonly identified by words such as "would," "may," "expects," "believes," "plans," "intends," "projects" and other terms with similar meaning. Investors are cautioned that the forward-looking statements in this document are based on current beliefs, assumptions and expectations, speak only as of the date of this press release and involve risks and uncertainties that could cause actual results to differ materially from current expectations. Such statements are subject to certain known and unknown risks and uncertainties, many of

which are difficult to predict and generally beyond AMD's control, that could cause actual results and other future events to differ materially from those expressed in, or implied or projected by, the forward-looking information and statements. Material factors that could cause actual results to differ materially from current expectations include, without limitation, the following: Intel Corporation's dominance of the microprocessor market and its aggressive business practices may limit AMD's ability to compete effectively; AMD is party to a wafer supply agreement with GLOBALFOUNDRIES Inc. (GF) with obligations to manufacture products at GF with certain exceptions. If GF is not able to satisfy AMD's manufacturing requirements, its business could be adversely impacted; AMD relies on third parties to manufacture its products, and if they are unable to do so on a timely basis in sufficient quantities and using competitive technologies, AMD's business could be materially adversely affected; failure to achieve expected manufacturing yields for AMD's products could negatively impact its financial results; the success of AMD's business is dependent upon its ability to introduce products on a timely basis with features and performance levels that provide value to its customers while supporting and coinciding with significant industry transitions; if AMD cannot generate sufficient revenue and operating cash flow or obtain external financing, it may face a cash shortfall and be unable to make all of its planned investments in research and development or other strategic investments; the loss of a significant customer may have a material adverse effect on AMD; AMD's receipt of revenue from its semi-custom SoC products is dependent upon its technology being designed into third-party products and the success of those products; global economic uncertainty may adversely impact AMD's business and operating results; the markets in which AMD's products are sold are highly competitive; AMD may not be able to generate sufficient cash to service its debt obligations or meet its working capital requirements; AMD has a substantial amount of indebtedness which could adversely affect its financial position and prevent it from implementing its strategy or fulfilling its contractual obligations; the agreements governing AMD's notes and the secured revolving line of credit (Secured Revolving Line of Credit) impose restrictions on AMD that may adversely affect its ability to operate its business; uncertainties involving the ordering and shipment of AMD's products could materially adversely affect it; the demand for AMD's products depends in part on the market conditions in the industries into which they are sold. Fluctuations in demand for AMD's products or a market decline in any of these industries could have a material adverse effect on its results of operations; AMD's ability to design and introduce new products in a timely manner is dependent upon third-party intellectual property; AMD depends on third-party companies for the design, manufacture and supply of motherboards, software and other computer platform components to support its business; if AMD loses Microsoft Corporation's support for its products or other software vendors do not design and develop software to run on AMD's products, its ability to sell its products could be materially adversely affected; and AMD's reliance on third-party distributors and Add-in-Board partners subjects it to certain risks. Investors are urged to review in detail the risks and uncertainties in AMD's Securities and Exchange Commission filings, including but not limited to AMD's Quarterly Report on Form 10-Q for the quarter ended September 24, 2016.

1 Data based on AMD Engineering design of Vega. Radeon R9 Fury X has 4 geometry engines and a peak of 4 polygons per clock. Vega is designed to handle up to 11 polygons per clock with 4 geometry engines. This represents an increase of 2.6x. VG-3

2 Discrete AMD Radeon™ and FirePro™ GPUs based on the Graphics Core Next architecture consist of multiple discrete execution engines known as a Compute Unit ("CU"). Each CU contains 64 shaders ("Stream Processors") working together. GD-78

Contact Information
Chris Hook
AMD Communications
512-578-9727
chris.hook@amd.com

Alina Ostrovsky
AMD Investor Relations
+1 (408) 749-6688
Alina.Ostrovsky@amd.com

Source: Advanced Micro Devices