November 16, 2015

AMD Launches 'Boltzmann Initiative' to Dramatically Reduce Barriers to GPU Computing on AMD FirePro(TM) Graphics

New Tools Target an Unprecedented 28 Teraflops of Processing at Less Than a Kilowatt by 2016

AUSTIN, TX -- (Marketwired) -- 11/16/15 -- Building on its strategic investments in heterogeneous system architecture (HSA), AMD (NASDAQ: AMD) announced a suite of tools designed to ease development of high-performance, energy efficient heterogeneous computing systems. The "Boltzmann Initiative" leverages HSA's ability to harness both central processing units (CPU) and AMD FirePro™ graphics processing units (GPU) for maximum compute efficiency through software. The first results of the initiative are featured this week at SC15 and include the Heterogeneous Compute Compiler (HCC); a headless Linux® driver and HSA runtime infrastructure for cluster-class, High Performance Computing (HPC); and the Heterogeneous-compute Interface for Portability (HIP) tool for porting CUDA-based applications to a common C++ programming model. The tools are designed to drive application performance across markets ranging from machine learning to molecular dynamics, and from oil and gas to visual effects and computer-generated imaging.

"AMD's Heterogeneous-compute Interface for Portability enables performance portability for the HPC community. The ability to take code that was written for one architecture and transfer it to another architecture without a negative impact on performance is extremely powerful," said Jim Belak, co-lead of the U.S. Department of Energy's Exascale Co-design Center in Extreme Materials and senior computational materials scientist at Lawrence Livermore National Laboratory. "The work AMD is doing to produce a high-performance compiler that sits below high-level programming models enables researchers to concentrate on solving problems and publishing groundbreaking research rather than worrying about hardware-specific optimizations."

New Compiler for Heterogeneous Computing

The promise of combining multi-core, serial processing CPUs with parallel-processing GPUs to maximize compute efficiency is already being seen in the industry, as driven by the Heterogeneous Systems Architecture (HSA) Foundation that counts AMD as a founding member. One of the goals for HSA is easing the development of parallel applications through use of higher level languages. The new AMD "Boltzmann Initiative" suite includes an HCC compiler for C++ development, greatly expanding the field of programmers who can leverage HSA. The new HCC C++ compiler is a key tool in enabling developers to easily and efficiently apply the hardware resources in heterogeneous systems. The compiler offers more simplified development via single source execution, with both the CPU and GPU code in the same file. The compiler automates the placement code that executes on both processing elements for maximum execution efficiency.

"Just as our customers are excited about our hardware innovation, including the introduction of the first GPU with High Bandwidth Memory this year and our new x86 core architecture coming next year, our innovations in software development are equally as important to them," said Mark Papermaster, senior vice president and chief technology officer, AMD. "The challenge has always been to unlock the hardware's capabilities and make them easily accessible to developers working to solve difficult problems. AMD's newest offering provides the keys to more readily access our parallel computing engines -- both multicore CPUs and GPUs -- and to make these benefits available to the mainstream of developers across a broad spectrum of computing platforms, from embedded to supercomputing."

Linux Driver and Runtime Focused on the Needs of HPC Cluster-Class Computing
To complement the new compilation tools, AMD has developed a new HPC-focused driver and system runtime. This new headless Linux driver brings key capabilities to address core high-performance computing needs, including low latency compute dispatch and PCIe® data transfers; peer-to-peer GPU support; Remote Direct Memory Access (RDMA) from InfiniBand™ that interconnects directly to GPU memory; and Large Single Memory Allocation support.

HIP-ifying CUDA Application to Run on AMD GPUs
To bring applications written for CUDA onto AMD platforms, AMD announces the new HIP tool. AMD testing shows that in many cases 90 percent or more of CUDA code can be automatically converted into C++ by HIP with the final 10 percent converted manually in the widely popular C++ language. This greatly expands the installed hardware base available to run what were formerly exclusively CUDA-based applications. At SC15, AMD is demonstrating the potential for HIP, running the CUDA-generated Rodinia benchmark suite on AMD GPUs.

Availability
An early access program for the "Boltzmann Initiative" tools is planned for Q1 2016.

Other News at SC15
AMD continues to support the open standard OpenCL™ programming language, giving developers low-level GPU programming and optimization capabilities. AMD has expanded its GPU compute libraries, allowing developers to easily accelerate compute-intensive functions by leveraging GPU compute. The addition of HcBLAS and HcFFT builds on work done with clBLAS, clFFT, clSPARSE and clRNG for OpenCL. The libraries include functions that are widely used in scientific computation and data science.

Industry Support for AMD
Support continues to grow for AMD's standards-based approach to HPC. AMD's investment in hardware and software development has resulted in a growing number of system and solution vendors that fully support AMD FirePro™ in enterprise environments, including ASRock, ASUS, Cirrascale, Cubix, Dell, Gigabyte, HP, One Stop Systems and Supermicro.

Demonstrations at SC15
At SC15 in booth 727, AMD will showcase:

- The new HSA-compliant C++ compiler, full open source HSA programming suite and HIP porting tools.
- Systems from Dell, HP and Supermicro that include AMD FirePro™ GPUs running demonstrations such as AMD FireRender, Abaqus and TUM Navier-Strokes.

- AMD's Opteron™ A1100 series ARM processor in SoftIron's new Enterprise Class Overdrive 3000 system for developers and Silver Lining System's ARM-Metal-as-a-Service (AMAAS) solution.
- The AMD FirePro™ S9170 server GPU in action with Dell and Supermicro servers, showcasing a variety of workloads that take advantage of the AMD FirePro S9170 card's 32GB high-speed and high-bandwidth onboard memory, making it ideal for accelerating large dataset computations.
- AMD engineers are scheduled to present two related tech talks
    - Tuesday, November 17 at 4:15 p.m. Greg Rodgers, design engineer, AMD, "Advances in HSA Software Stack"
    - Wednesday, November 18 at 5:15 p.m. Srikanth Gollapudi, PMTS software engineer, AMD, "Porting CUDA applications to C++ with New AMD HIP conversion tools"

Supporting Resources

- Blog:
  A Defining Moment for Heterogeneous Computing by AMD CTO Mark Papermaster
- AMD Developer Central
- Learn more about Heterogeneous System Architecture
- Become a fan of AMD on Facebook
- Follow AMD on Twitter @AMD and @AMDFirePro

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, FirePro, Opteron, and combinations thereof, are trademarks of Advanced Micro Devices, Inc. OpenCL is a trademark of Apple Inc. used by permission by Khronos. Other names are for informational purposes only and may be trademarks of their respective owners.

Contact:
Gary Silcott
AMD Communications
(510) 602-0889
Email Contact

John Swinimer
AMD Communications
(289) 695-0600
Email Contact

Source: Advanced Micro Devices