

The background of the slide is a high-magnification, colorful microchip. The chip is divided into various sections, with colors ranging from bright orange and red to yellow and green. The intricate patterns of the chip's circuitry are visible. A large, solid purple shape is overlaid on the bottom left and bottom center of the slide, partially obscuring the chip image. A green arrow points from the top left towards the center, and another green arrow points from the right towards the end of the main title.

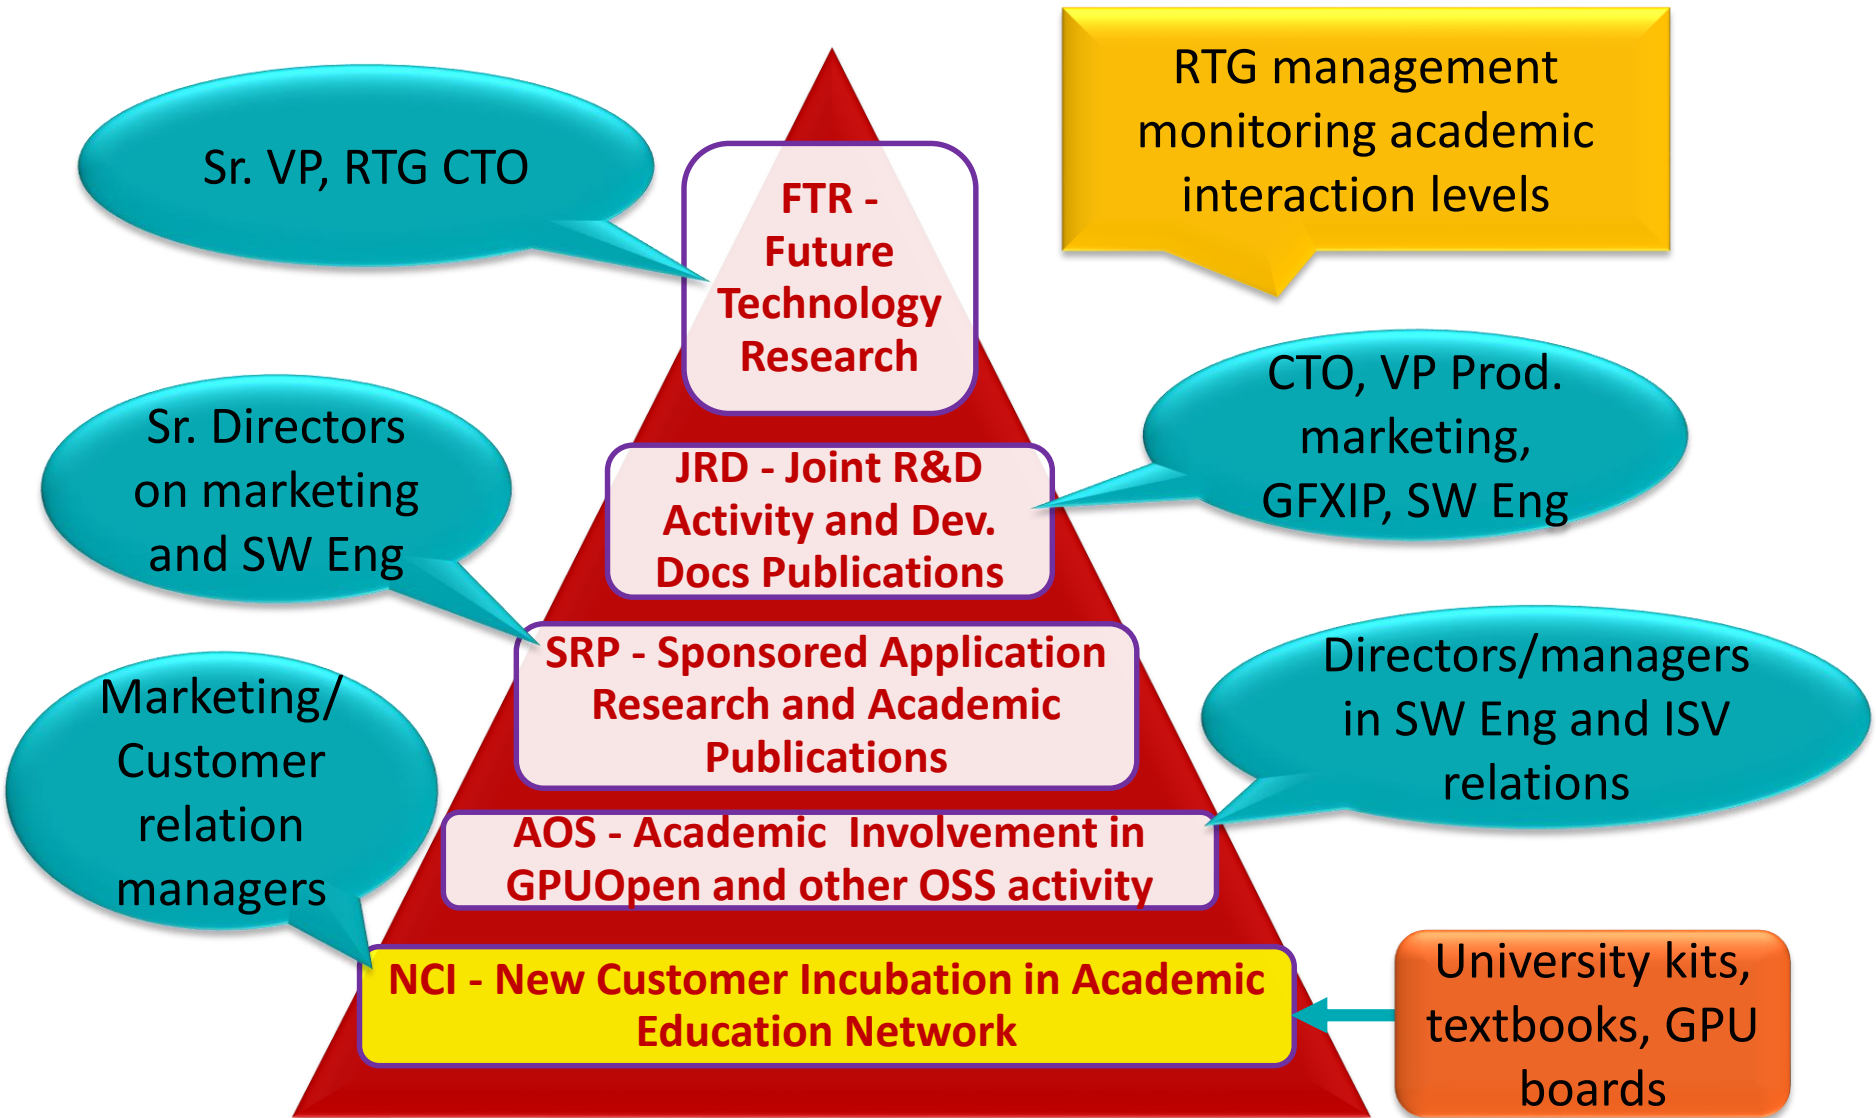
RADEON TECHNOLOGY GROUP GLOBAL ACADEMIC CONNECTIONS

DR. TIMOUR PALTASHEV,
RADEON TECHNOLOGY GROUP
AUGUST 15, 2016

HIERARCHY OF ACADEMIC RELATIONS



INTERACTION LEVELS AND MANAGEMENT RESPONSIBILITY



AMD/RTG ACADEMIC RELATIONS LEVELS 1



INTRODUCTORY LEVEL OF ACADEMIC AND DEVELOPER RELATIONS

- ▲ **NCI** - New Customer Incubation in Academic Education Network
 - Massive incursion of AMD technology and product in global college education network (similar to NVidia and Intel)
 - Assigned manager to deal with multiple institutions globally
 - Textbooks, university kits (partially existing) and remote cluster access or GPU board support for teaching/course projects
- ▲ **AOS** - Academic Involvement in GPUOpen and other OSS activity
 - Attracts a lot of interests from academia but needs well-organized technical documentation and consultancy support on RTG side
 - Could be multiple academic contributors for base libraries and certain application domain support packages (similar to NVidia academic network)
 - RTG ISV support should be arranged to work with academic teams using proper incentives and equipment access for OSS contribution

AMD/RTG ACADEMIC RELATIONS LEVELS 2



ADVANCED LEVELS OF ACADEMIC RELATIONS WITH RTG AMD

- ▲ **SRP** - Sponsored Application Research and Academic Publications
 - Selected academic institutions to be sponsored on research and publication activity to promote AMD/RTG technology
 - Proper incentives and awards targeted to cover existing gaps in certain application domains critical for RTG product sales increase
- ▲ **JRD** - Joint R&D Activity and Publications (typical for AMD Research)
 - Joint SDK tools and base libraries development/improvements with selected best academic entities for hot R&D topics defined by DoE, virtual reality, computer vision and DNN applications
 - Following academic publications for technology promotion
- ▲ **FTR** - Future Technology Research Collaboration
 - Very few most advanced academic entities to explore future technology possible implementations (probably in VR and HPC domains)

NCI LEVEL SUPPORT WITH UNIVERSITY PACKAGES



EASY WAY TO START NEW COURSES IN ACADEMIA

- ▲ **NCI - New Customer Incubation in Academic Education Network**
 - Corporate resources to form packages with university kits, textbooks and GPU boards to support college teaching & course projects
- ▲ **Compute Package with OpenCL 2.0 textbook, university kit and GPU board**
 - Funding to buy books and boards with following shipping to participating universities
 - Can be complemented with university kit based on new textbook “Heterogeneous System Architecture: A new compute platform infrastructure” and HCC materials from GPUOpen
- ▲ **Graphics Package with “OpenGL Superbible”, university kit and GPU board**
 - Needs resource supported lecture slides development with following funding to buy a number of books
 - Can be complemented with university kit based on new text “Vulkan Programming Guide: The Official Guide to Learning Vulkan (OpenGL)”

POWERFUL GRAPHICS ACCELERATOR FOR PARTNERS



RADEON R9 FURY: FIRST HIGH-VOLUME INTERPOSER WITH HBM MEMORY




- ▲ **8.5 Teraflops** on 1011 sq. mm of silicon interposer (with HBM memory)
- ▲ Graphics Core Next Architecture with 64 Compute Units and 4096 Stream Processors
- ▲ 596 sq. mm. Graphics Engine
- ▲ Targeted to Virtual Reality application with LiquidVR API
 - LiquidVR™ is an AMD initiative dedicated to making VR as comfortable and realistic as possible

COMPUTER GRAPHICS INTRODUCTION: OPENGL 4.5



POPULAR TEXTBOOK TO BE USED FOR BOTH INTRODUCTORY AND ADVANCED COURSES

- ▲ A practical introduction to real-time 3D graphics, including foundational math
- ▲ Core techniques for rendering, transformations, and texturing
- ▲ Shaders and the OpenGL Shading Language (GLSL) in depth
- ▲ Vertex processing, drawing commands, primitives, fragments, and framebuffers
- ▲ Compute shaders: harnessing graphics cards for more than graphics
- ▲ Pipeline monitoring and control
- ▲ Managing, loading, and arbitrating access to data
- ▲ Building larger applications and deploying them across platforms
- ▲ Advanced rendering: light simulation, artistic and non-photorealistic effects, and more
- ▲ Supercharging performance with persistent maps, bindless textures, and fine-grained synchronization




Any laptop with AMD A-series APU is perfectly fine to work with OpenGL 4.5

NEXT STEP AFTER OPENGL: VULKAN GRAPHICS API



NEW TEXTBOOK TO SUPPORT ACADEMIC PARTNERS ACTIVITY IN GRAPHICS

- ▲ Extensively tested code examples to demonstrate Vulkan's capabilities and show how it differs from OpenGL
- ▲ Expert guidance on getting started and working with Vulkan's new memory system
- ▲ Thorough discussion of queues, commands, moving data, and presentation
- ▲ Full explanations of the SPIR-V binary shading language and compute/graphics pipelines
- ▲ Detailed discussions of drawing commands, geometry and fragment processing, synchronization primitives, and reading Vulkan data into applications
- ▲ A complete case study application: deferred rendering using complex multi-pass architecture and multiple processing queues




Powerful Radeon R9 Fury X graphics accelerator is a part of university package

HETEROGENEOUS COMPUTING WITH OPENCL 2.0



SIGNIFICANT UPGRADE OF OPENCL 1.2 WITH ADVANCED CONCEPTS OF OPENCL 2.0

- ▲ New revision of OpenCL textbook for the latest OpenCL standard
 - Memory management (SVM)
 - Dynamic parallelism
 - Mapping of models onto latest device architectures
- ▲ Developed collaboratively between AMD and Northeastern University
 - Material suitable for industry developers or advanced university students
 - Strong focus on CPU, GPU, and APU architectures
- ▲ Compliant with Heterogeneous System Architecture Concept (HSA)




Works on any AMD platform even without graphics accelerator
OpenCL 1.2 works on any platform

NEW RESEARCH FIELD IN PARALLEL COMPUTING PLATFORM **AMD**

HETEROGENEOUS SYSTEM ARCHITECTURE: A NEW COMPUTE PLATFORM INFRASTRUCTURE

- ▲ Provides clear and concise explanations of key HSA concepts and fundamentals by expert HSA Specification contributors
- ▲ Explains how performance-bound programming algorithms and application types can be significantly optimized by utilizing HSA hardware and software features
- ▲ Presents HSA simply, clearly, and concisely without reading the detailed HSA Specification documents
- ▲ Demonstrates ideal mapping of processing resources from CPUs to many other heterogeneous processors that comply with HSA Specifications
- ▲ **This field is very promising for R&D in academia**



AMD may provide special boxes or laptops with enabled HSA features for research

OPENCL 2.0 UNIVERSITY KIT



- ▲ Course materials developed to accompany Heterogeneous Computing with OpenCL 2.0 textbook
 - Designed to facilitate adoption into university programs
- ▲ Materials progress along with textbook chapters
- ▲ Includes source code for examples presented in the textbook

<http://developer.amd.com/partners/university-programs/opencl-university-kits/>



OPENCL UNIVERSITY PROGRAMS



- ▲ AMD OpenCL University Programs website encompasses OpenCL initiatives between AMD and academia
 - Textbook
 - University kit with per-chapter lecture slide sets
 - Active university programs
 - Possible remote access to clusters
 - Open-source releases (more coming soon)
- ▲ <http://developer.amd.com/partners/university-programs/>



Heavy computational workloads have traditionally been processed on a CPU, but technology is shifting to a new computing paradigm that relies more on the GPU or an APU. OpenCL™ is the widely adopted industry standard for running parallel tasks on CPUs, APUs and GPUs using the same code. Check out some of the university resources below.

AOS - ACADEMIC OPEN-SOURCE STRATEGY



SOLVE OUR OSS AND BASE LIBRARIES PROBLEMS WITH SIGNIFICANT INVOLVEMENT OF ACADEMIC PARTNERS

- ▲ Base libraries domain: BLAS, SPARSE, RAND, FFT, Convolution etc.
 - GCN++ friendly code refactoring and performance tune-up
 - Joint library development
- ▲ DNN domain to be extensively developed
- ▲ CUDA → HIP → C++ → HCC support may require:
 - Replication of popular NV API and template libraries in our OSS SW stack: Thrust, CUB
 - NV ModernGPU library concept implementation on our code base and GCN architecture
- ▲ Marketing and Application engineering managers define hot topics and budget resources for academic incentives
 - RTG CTO people can help to select trustable academic entities and arrange cooperation flow with RTG SW teams

GPUOPEN INITIATIVE



BOTH GRAPHICS AND COMPUTE API: OPENGL/VULKAN AND OPENCL/HSA/ROCM

- ▲ <http://gpuopen.com/> AMD initiative designed to enable developers to create ground-breaking PC games, computer generated imagery and GPU computing applications for great performance and lifelike experiences using no cost and open development tools and software
- ▲ <http://gpuopen.com/games-cgi/> GPUOpen Games & CGI is designed to empower graphics developers to create ground-breaking games and computer-generated imagery with great performance and vivid visual effects using open source software and tools
- ▲ <http://gpuopen.com/professional-compute/> GPUOpen Professional Compute is designed to empower all types of developers to accelerate the implementation of their vision and help solve their biggest challenges in instinctive and high-performance GPU computing through optimized open-source driver/runtimes and standard-based languages, libraries and applications
- ▲ <https://github.com/RadeonOpenCompute/ROCm> Open Heterogenous Computing Platform (Linux(R) Driver and Runtime Stack) optimized for HPC & Ultra-scale class computing (heterogeneous C and C++ Single Source)

WHEN 'OPEN' REALLY MEANS OPEN



▲ AMD OpenCL open source libraries

- Source code available on GitHub
- Developers have full control and code ownership

▲ AMD OpenCL libraries run on GPU and CPU

- Tuned by AMD for maximum efficiency on AMD GPUs



cBLAS

The primary goal of cBLAS is to make it easier for developers to utilize the inherent performance and power efficiency benefits of heterogeneous computing. cBLAS interfaces do not hide nor wrap OpenCL interfaces, but rather leaves OpenCL™ state management to the control of the user to allow for maximum performance and flexibility.



cSparse

True in spirit with the other clMath libraries, cSparse exports a “C” interface to allow projects to build wrappers around cSPARSE in any language they need. cSPARSE is an OpenCL™ library implementing Sparse linear algebra. This project is a result of a collaboration between AMD Inc. and Vrtis Ltd.



cRNG

cRNG has both host and device side interfaces to give you, the developer, sufficient control. The host side interfaces are native C APIs. cRNG 1.0 is an OpenCL library that generates uniform random numbers.



cFFT

cFFT is a software library containing Fast Fourier Transform (FFT) functions written in OpenCL. In addition to GPU devices, the libraries also support running on CPU devices to facilitate debugging and heterogeneous programming.

AMD open source OpenCL libraries : <https://github.com/clMathLibraries>

OpenCL™ 2.0 is Here!

Learn more and tap into the tremendous performance potential of modern heterogeneous systems with the latest AMD APP SDK.

[Read the blog.](#)

OpenCL™ 2.0 Samples



Demystifying OpenCL™ 2.0 features

The OpenCL™ 2.0 samples posted here accompany a series of blog posts where we attempt to demystify the most important OpenCL™ 2.0 features by investigating their significance as well as how and under what conditions they improve programmability and performance. The best way to learn about the features, of course, is to use them:

- Download and install the AMD OpenCL 2.0 Driver located [here](#)
- Download the code sample(s) from the table below and run them on any of the AMD platforms listed on the driver download page (link above). See the Readme file in the sample package for further details.

Downloads

| File Name | Version | Size | Launch Date | OS | Bitness | Description |
|---|---------|--------|-------------|-------------------|---------|---|
| Binary Search using OpenCL 2.0 SVM | | | | | | |
| SVMBinarySearchTree.zip | N/A | 124 KB | 10/24/2014 | Windows and Linux | 64-bit | Sample code showing use of OpenCL 2.0 SVM feature in a binary tree search |
| Producer/consumer sample using OpenCL 2.0 Pipes | | | | | | |
| PipeProducerConsumerKernels.zip | N/A | 128 KB | 10/31/2014 | Windows and Linux | 64-bit | Sample code implementing a producer/consumer construct using OpenCL 2.0 pipes |



AMD Accelerated Parallel Processing OpenCL™ User Guide

- ☰ Preface
- ☰ Contents
- ☰ Chapter 1 OpenCL Architecture and AMD Accelerated Parallel Processing
- ☰ Chapter 2 AMD Implementation
- ☰ Chapter 3 Building and Running OpenCL Programs
- ☰ Chapter 4 Debugging and Profiling OpenCL
- ☰ Chapter 5 OpenCL Static C++ Programming Language
- ☰ Chapter 6 OpenCL 2.0
- ☰ Appendix A OpenCL Optional Extensions
- ☰ Appendix B The OpenCL Installable Client Driver (ICD)
- ☰ Appendix C OpenCL Binary Image Format (BIF) v2.0
- ☰ Appendix D Hardware overview of pre-GCN devices
- ☰ Appendix E OpenCL-OpenGL Interoperability
- ☰ Appendix F New and deprecated functions in OpenCL 2.0
- ☰ Appendix G Standard Portable Intermediate Representation (SPIR)
- ☰ Index

- ▲ Primary requirements for HSA-compliant system defined by “**HSA Platform System Architecture Specification Version 1.0**”

http://www.hsafoundation.com/html/HSA_Library.htm

1. **A shared virtual memory system satisfying the HSA Memory Consistency Model**
2. **User mode queues providing a standard interface to the system’s agents**
3. **Support for AQL Queuing Language**
 - AQL – Architected Queuing Language which defines interaction between agents in HSA system
4. **Heterogeneous System Architecture Intermediate Language (HSAIL) for agent’s code (kernels)**

Other important documents (same link):

- ▲ **HSA Programmer's Reference Manual Version 1.0**
- ▲ **HSA Runtime Programmer’s Reference Manual Version 1.0**

CONTACT INFORMATION



▲ Contacts for OpenCL-related academic initiatives

▲ Timour Paltashev

– Senior manager, Radeon Technology Group, GPU architecture and global academic connections

– timour.paltashev@amd.com

▲ Dana Schaa

– GPU architecture, textbook author, university kit contributor

– dana.schaa@amd.com



QUESTIONS AND ANSWERS

