



# AMD University Program

## Research & Advanced Development

**Andrew Schmidt**

AMD University Program – North America

[Andrew.Schmidt@amd.com](mailto:Andrew.Schmidt@amd.com)

**AMD**   
together we advance\_

---

# Agenda

- 
1. AMD University Program Overview
  2. Heterogenous Accelerated Compute Clusters (HACC)
  3. AMD AI & HPC Cluster Program
  4. Questions

# AMD University Program Vision

Empower academics with AMD technology to enhance teaching and learning experiences and advance state-of-the-art research.




[www.amd.com/AUP](http://www.amd.com/AUP)

# AMD University Program Team

Dedicated world-wide technical team

Supporting High Performance and Adaptive Compute

25+ years experience working with academia



# Academic and Research Community



## Students

Affordable access  
to AMD products

Build industry relevant skills and  
knowledge with AMD products



## Educators

University resources  
and support

Deliver exceptional teaching with  
AMD educational resources and support




## Researchers

Leading-edge  
AMD technologies

Advance state-of-the-art  
research

# AMD University Team Offerings



**Research Programs**



**Donation Program**



**Teaching Resources**



**Training**



**Academic Solutions**




**Support**

# Donation Program

Request academic and AMD boards for teaching and research

Apply now for academic licenses for AMD Software & IP


Request silicon devices for custom projects




**Development boards**



**Production boards**



**Academic boards**



**Silicon devices**

[www.amd.com/AUP](http://www.amd.com/AUP)

# Teaching Resources

Open-source teaching material

Video presentations

Tech Talks

Tutorial Labs

*Academic Experts* program  
for access to advance training



[www.amd.com/AUP](http://www.amd.com/AUP)



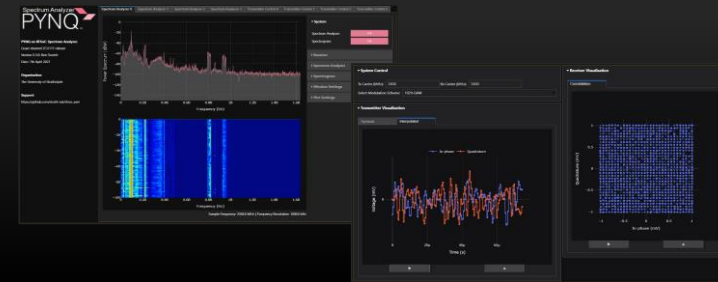


# Enabling Hands-on Learning and Exploration


## AMD Text Books



## Jupyter Notebooks




## ZCU111 and RFSoc 4x2



## Open-source framework



# University Training



Live and remote  
instructor-led  
training

PhD schools

Free to attend  
for academics

Hands-on lab  
exercises

Supported on University  
Program hardware and  
cloud resources

[www.amd.com/AUP](http://www.amd.com/AUP)

# AMD Partners

## University Program

Leads interaction with university community

Coordinates collaborations across AMD partners

Facilitates research opportunities



# RESEARCH PROGRAMS

## Accelerating AI and Science

**AMD**   
**AI & HPC Fund**



# Heterogeneous Accelerated Compute Clusters

Enabling Novel Research in Heterogeneous  
Compute Acceleration for HPC

# Heterogenous Accelerated Compute Clusters (HACCs)


Supports novel research in heterogenous compute acceleration for AI and HPC to foster a community of leading academic teams to conduct state-of-the-art research

Encompasses systems, architecture, tools and applications

Established at six of world's most prestigious universities

Equipped with the latest AMD hardware and software technologies

## RESEARCH AREAS



# HACCs: Heterogeneous Accelerated Compute Clusters

Remote access to  
Adaptive Compute hardware

HACC user group meetings

Access to AMD researchers

Collaboration opportunities

  
UCLA ILLINOIS

 ETH zürich

 Indian  
Institute  
of Science

 NUS  
National University  
of Singapore

 AMD  
EPYC

 AMD  
INSTINCT


 AMD  
ALVEO

 AMD  
VERSAL

[www.amd-haccs.io](http://www.amd-haccs.io)

★ Newest HACC at IISc, Bangalore

# HACC Adaptive Computing Hardware




- HACC hardware consists of:
  - Compute and Alveo™ nodes (initially U250 and U280 with HBM)
  - Latest heterogeneous nodes (SMC 4124GS) include:
    - 2 EPYC™ 3rd generation CPUs
    - 4 AMD Instinct™ MI210 GPUs
    - 2 Alveo U55C FPGA with HBM
    - 2 VCK-5000 ACAP/Versal with AIEs
    - Run-time via AMD ROCm™, XRT
    - SW development via HIP, Vitis, frameworks
  - 100G network
- Community hub for researchers
  - Support from in-house AMD research groups
  - Reproducible results & experiments





# HACC Resources



## HACC

- HACC Home
- Get Started
- Acknowledgment
- Research Awards


## Centers

- ETH Zurich
- IISc
- NUS
- Paderborn
- UCLA
- UIUC

## Resources

- Events
- Publications
- Examples
- Frameworks
- Tools
- Training

## Heterogeneous Accelerated Compute Clusters



The Heterogeneous Accelerated Compute Clusters (HACC) program is a special initiative to support novel research in adaptive compute acceleration for high performance computing (HPC). The scope of the program is broad and encompasses systems, architecture, tools and applications.

Five HACCs have been established at some of world's most prestigious universities. HACCs are equipped with the latest AMD Adaptive Computing hardware and software technologies. Each cluster is specially configured to enable some of the world's foremost academic teams to conduct state-of-the-art HPC research.

See the [getting started](#) page for more details on the program and how to apply.

<https://www.amd-haccs.io/>

# HACC Tech Talk Series

## HACC Tech Talk series

The talks will be hosted as Zoom webinars and are free to attend. The format for each session will be two 30-minute talks on topics related to Adaptive Compute.

If you would like to present your work at an HACC Tech Talk, please contact [ap@amd.com](mailto:ap@amd.com) with an outline of your proposal.

## Past talks

### HACC Tech Talk 15

14th November 2024

#### Optimizing Communication for Latency Sensitive HPC Applications on up to 48 FPGAs Using ACCL

- Send data from buffer in global memory of remote FPGA
- Received data is stored in dedicated receive buffers in global memory of FPGA
- Copied to final destination during receive call (Blocking the CCLD)
- Communication latency can be modelled with

Home Videos Playlists

OpenMP Offloading 34:15

AMD HACC Tech Talks: OMPC FPGA: Integrating Multi-FPGA Acceleration to... 38 views • 2 days ago

Buffered Communication 28:09

AMD HACC Tech Talks: Optimize Latency Sensitive HPC Application Communication... 12 views • 2 days ago

How does the workload size/shape affect kernel efficiency? 26:38

AMD HACC Tech Talks: Performance Analysis of GEMM Workloads on the AMD... 29 views • 2 days ago

Electronic-skin Interface 19:52

AMD HACC Tech Talks: Energy-Efficient Computation for Machine Learning... 33 views • 2 days ago

Programming the AI Engine 23:09

BALBOA: RoCE-Stack from Coyote 29:52

Top-Level Architecture 25:52

Pruning: Structured versus Unstructured 28:02

AMD HACC Tech Talks: TaPaSCo on Versal: AIE, QDMA Streaming and MRMAC 89 views • 4 weeks ago

AMD HACC Tech Talk: RoCE Balboa – RDMA Deep Packet Inspection at Line Rate with... 96 views • 4 weeks ago

AMD HACC Tech Talk: Energy-efficient 2.5D System for LLM Inference 164 views • 1 month ago

AMD HACC Tech Talk: Pruning in the Datafree regime, A Distributional approach... 66 views • 1 month ago

Let's use Rocgdb to debug our Ghost\_Exchange hip code 30:18

A Tale of Host and Device 33:13

Implementations: Single SpMM pipeline (1 Gflop) 31:24

Sparse Processing Architecture: Off/On-Chip Memory 29:41

AMD HACC Tech Talk: Profiling Tools for 30:18

AMD HACC Tech Talk: ROCm Ecosystem and 33:13

AMD HACC Tech Talk: HiHiSPMV: with 31:24

AMD HACC Tech Talk: Serpens and 29:41

[www.amd-haccs.io/hacc\\_tech\\_talks.html](http://www.amd-haccs.io/hacc_tech_talks.html)

**Next Talk December 9th!**

# Join the HACC Program

- Researchers working in relevant areas are invited to apply to join the HACC.
- Partner research teams can remotely access the clusters' computing resources to carry out research
- HACCs are a community hub for researchers to come together to collaborate with other experts in the field, including AMD in-house research groups.
- Visit the HACC resources pages to find out more about the centers, how to get started, events as well as related publications, examples, frameworks, tools and training material.
- Please contact [research\\_clusters@amd.com](mailto:research_clusters@amd.com) with any questions related to this program.

**<https://www.amd-haccs.io/get-started.html>**

A woman with her hair in a ponytail is seen from the side, working on a laptop in a server room. The room is filled with rows of server racks, and the lighting is dim, with a blue tint. The text is overlaid on the center of the image.

# **AMD AI & HPC Cluster Program**

Accelerating research and education in AI and HPC

# 20+

**PETAFLUPS**

COMPUTING POWER

combined that would  
rank among the **fastest**  
supercomputers in the world\*

## AMD AI & HPC Fund

\*According to the most recent Top 500 list

AMD UNIVERSITY PROGRAM

# \$31M

TOTAL  
MARKET  
VALUE

# 28 ▶ 9

GRANTEES

COUNTRIES

US • CANADA • FRANCE • GERMANY  
ITALY • SWITZERLAND • INDIA • SINGAPORE

# 400+

CPUs

AMD   
EPYC

# 2800+

GPUs


AMD   
INSTINCT

# 100+

Alveo Cards

AMD   
ALVEO

# AMD AI & HPC Cluster



40-node cluster consisting of AMD EPYC CPUs and AMD Instinct GPUs (MI100, MI210, MI250, MI300X) connected with a high-speed GPU-aware interconnect

- ROCm platform w/ HIP for explicit GPU programming
- ROCm-enabled PyTorch and TensorFlow for AI workloads
- Slurm for job management
- Singularity for container solutions
- Persistent file storage on parallel file system(s)




Allocations available to academic researchers and educators through our external user program




# AMD AI & HPC Cluster Program

Perform Your Academic Research on Our 40-Node AMD EPYC + Instinct (MI210, MI250, MI300X) Cluster


**MI210 Node**



**MI250 Node**



**MI300X Node**



**Proposal Form**



# AMD AI & HPC Cluster User Program – Proposal Requirements

## Project PIs

- Must be researchers or educators from academic organizations (e.g., universities, national labs, technical institutes)
- Must be postdoc or higher



## Project Proposals

- Proposed work must be non-commercial
- Proposed work must be for 1 year or less (follow on allocations require new proposal)
- Node-hour requests must be appropriate size for cluster (typical allocations are between 5K – 15K node-hours)
  - Total node-hours per year = (39 compute nodes) \* (365 days) \* (24 hours / day) = 341,640 node-hours
- Project outcome should be highly impactful (e.g., publications, open-source libraries, educational content)

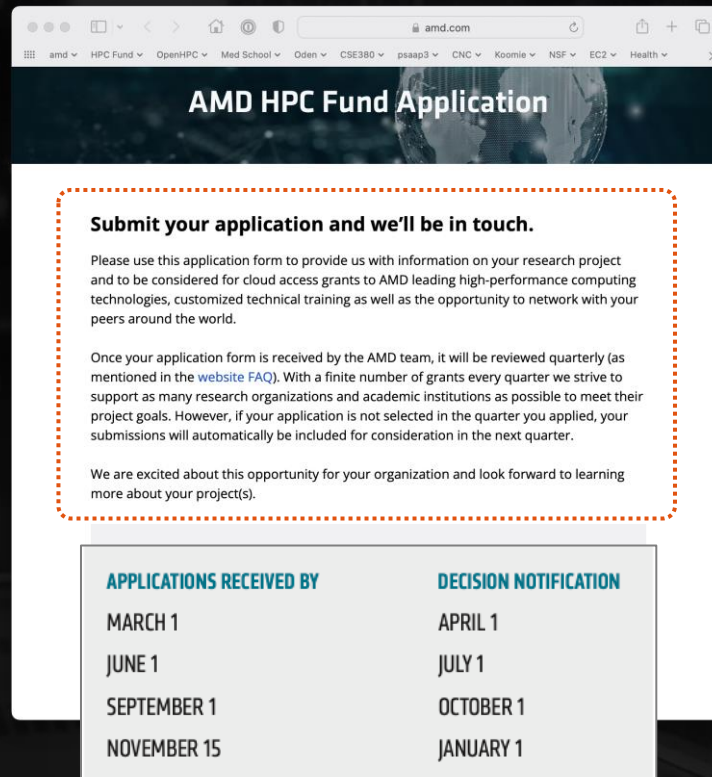


# AMD AI & HPC Cluster User Program – How to Request Access

## Proposal Form

<https://www.amd.com/en/forms/registration/amd-hpc-fund-research-accelerator.html>

*Research applications in AI/ML are particularly encouraged!*



The screenshot shows the AMD HPC Fund Application page. The header includes the AMD logo and navigation links for various departments like HPC Fund, OpenHPC, Med School, Oden, CSE380, psaaap3, CNC, Koomie, NSF, EC2, and Health. The main heading is "AMD HPC Fund Application". Below this, a dashed orange box contains the following text:

**Submit your application and we'll be in touch.**


Please use this application form to provide us with information on your research project and to be considered for cloud access grants to AMD leading high-performance computing technologies, customized technical training as well as the opportunity to network with your peers around the world.

Once your application form is received by the AMD team, it will be reviewed quarterly (as mentioned in the [website FAQ](#)). With a finite number of grants every quarter we strive to support as many research organizations and academic institutions as possible to meet their project goals. However, if your application is not selected in the quarter you applied, your submissions will automatically be included for consideration in the next quarter.

We are excited about this opportunity for your organization and look forward to learning more about your project(s).

APPLICATIONS RECEIVED BY	DECISION NOTIFICATION
MARCH 1	APRIL 1
JUNE 1	JULY 1
SEPTEMBER 1	OCTOBER 1
NOVEMBER 15	JANUARY 1





# FUNDING TRANSFORMATIONAL TECHNOLOGY AND ENGINEERING



Institutional Research

# Funding Academic Research Program Overview

“Accelerate the world’s transformative technology and engineering”



Enables groundbreaking and innovative research projects through grants



Proposals for grants are submitted through AMD University Program and will be evaluated and selected by a cross-functional team

[www.amd.com/AUP](http://www.amd.com/AUP)

# Contact Us

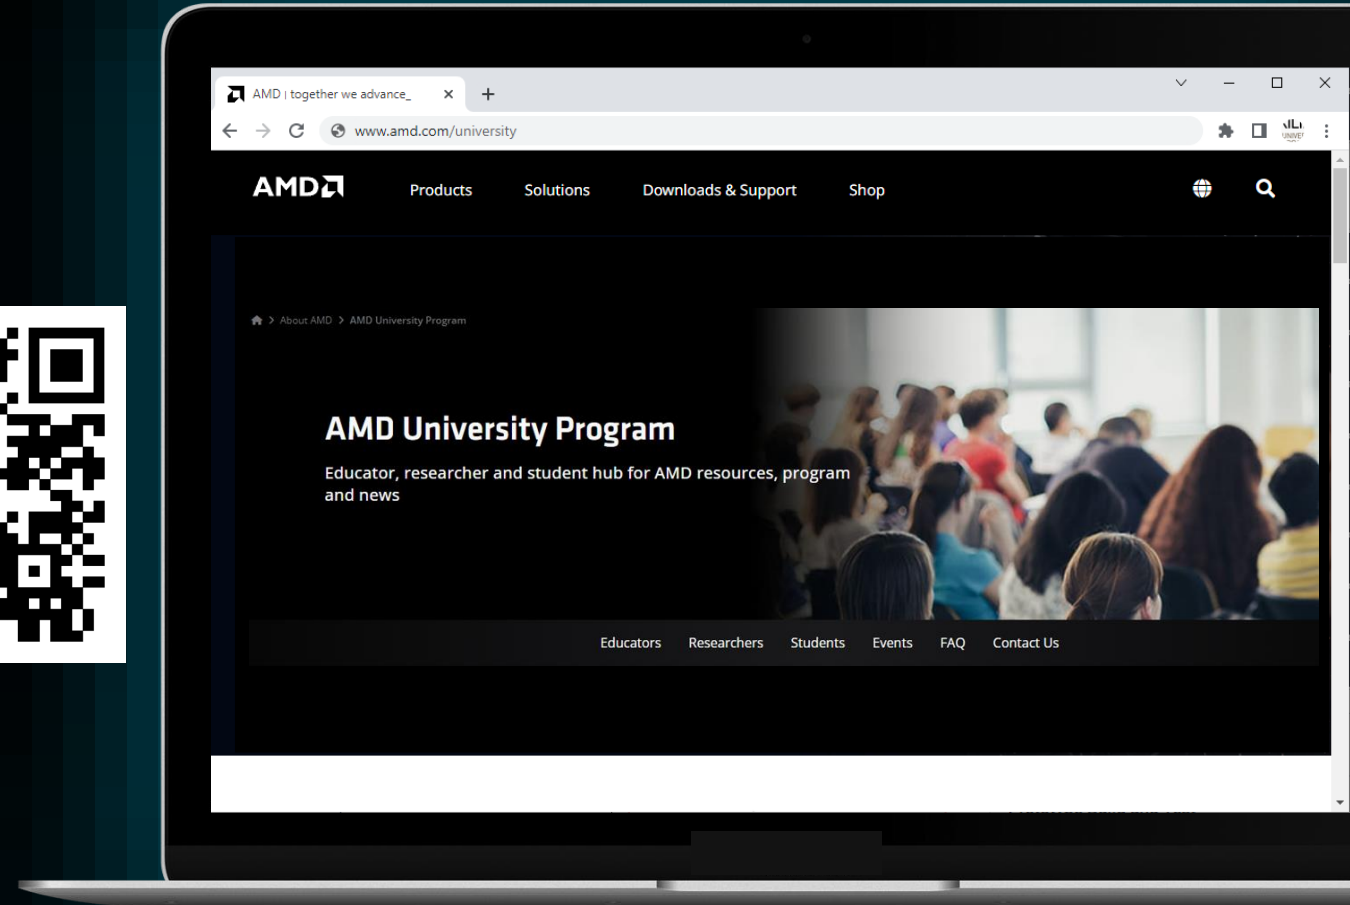
## Visit our website to:

- Discover our research programs
- Access educational resources
- Submit a donation request
- Find training & other events



## Email us:

[aup@amd.com](mailto:aup@amd.com)



[www.amd.com/AUP](http://www.amd.com/AUP)

# DISCLAIMER

The information presented in this document is for informational purposes only and may contain technical inaccuracies, omissions, and typographical errors. The information contained herein is subject to change and may be rendered inaccurate for many reasons, including but not limited to product and roadmap changes, component and motherboard version changes, new model and/or product releases, product differences between differing manufacturers, software changes, BIOS flashes, firmware upgrades, or the like. Any computer system has risks of security vulnerabilities that cannot be completely prevented or mitigated. AMD assumes no obligation to update or otherwise correct or revise this information. However, AMD reserves the right to revise this information and to make changes from time to time to the content hereof without obligation of AMD to notify any person of such revisions or changes. THIS INFORMATION IS PROVIDED 'AS IS.' AMD MAKES NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE CONTENTS HEREOF AND ASSUMES NO RESPONSIBILITY FOR ANY INACCURACIES, ERRORS, OR OMISSIONS THAT MAY APPEAR IN THIS INFORMATION. AMD SPECIFICALLY DISCLAIMS ANY IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY, OR FITNESS FOR ANY PARTICULAR PURPOSE. IN NO EVENT WILL AMD BE LIABLE TO ANY PERSON FOR ANY RELIANCE, DIRECT, INDIRECT, SPECIAL, OR OTHER CONSEQUENTIAL DAMAGES ARISING FROM THE USE OF ANY INFORMATION CONTAINED HEREIN, EVEN IF AMD IS EXPRESSLY ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

© 2024 Advanced Micro Devices, Inc. All rights reserved. AMD, the AMD Arrow logo, Alveo, AMD EPYC, AMD Instinct, AMD ROCm, Zynq, and combinations thereof are trademarks of Advanced Micro Devices, Inc. Other product names used in this publication are for identification purposes only and may be trademarks of their respective companies

**AMD** 

**together we advance\_**