

# Introduction to Programming AMD GPUs w/ HIP

## Presenter Bio

Tom Papatheodore (AMD University Program) is a senior member of technical staff (SMTS) working within AMD's University Program as a software development engineer. In this role, he helps to enable artificial intelligence (AI) and high-performance computing (HPC) research at universities and other academic institutions through donations of AMD hardware and/or allocations of time on AMD's academic AI & HPC Cluster.

## Agenda (4-Hour Tutorial)

Duration (minutes)	Topic
15	Access to AMD AI & HPC Cluster
15	Introduction of Tutorial & Motivation <ul style="list-style-type: none"><li>- GPU usage in HPC modeling &amp; simulation and AI</li><li>- AMD Instinct and Radeon GPUs</li></ul>
60	Introduction to ROCm/HIP (Part 1) <ul style="list-style-type: none"><li>- Overview of ROCm ecosystem and HIP</li><li>- GPU host-device programming model</li><li>- HIP API, kernel language</li><li>- Learn by examples w/ presenter demos<ul style="list-style-type: none"><li>- The basics w/ vector addition, error checking, GPU timers</li></ul></li></ul>
30	Hands-On Session (Part 1) <ul style="list-style-type: none"><li>- Exercises base on Part1 lecture/demos</li></ul>
60	Introduction to ROCm/HIP (Part 2) <ul style="list-style-type: none"><li>- Learn by examples w/ presenter demos<ul style="list-style-type: none"><li>- Concurrent kernels, overlap data transfers w/ compute, hipify</li></ul></li></ul>
55	Hands-On Session (Part 2) <ul style="list-style-type: none"><li>- Exercises based on Part 1 and 2 lecture/demos</li></ul>
5	Wrap-Up