



AMD Journey to Exascale


Mathieu Gontier

Field Application Engineer
AMD HPC Center of Excellence

HPCWS2021 (ECMWF)

September 22, 2021

Twitter: @mathieugontier



Weather forecast simulation
needs for exascale systems

Last revision: 15 September 2021



DISCLAIMER AND TRADEMARKS


DISCLAIMER The information contained herein is for informational purposes only, and is subject to change without notice. While every precaution has been taken in the preparation of this document, it may contain technical inaccuracies, omissions and typographical errors, and AMD is under no obligation to update or otherwise correct this information. Advanced Micro Devices, Inc. makes no representations or warranties with respect to the accuracy or completeness of the contents of this document, and assumes no liability of any kind, including the implied warranties of noninfringement, merchantability or fitness for particular purposes, with respect to the operation or use of AMD hardware, software or other products described herein. No license, including implied or arising by estoppel, to any intellectual property rights is granted by this document. Terms and limitations applicable to the purchase or use of AMD's products are as set forth in a signed agreement between the parties or in AMD's Standard Terms and Conditions of Sale. AMD, the AMD Arrow logo, EPYC™, Instinct™, Ryzen™, Ryzen™ Threadripper™, ROCm™ 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. Linux® is the registered trademark of Linus Torvalds in the U.S. and other countries. Windows® is a registered trademark of Microsoft Corporation in the US and/or other countries. LLVM™ is a trademark of LLVM Foundation. Java® is a registered mark of Oracle and/or its affiliates. PyTorch, the PyTorch logo and any related marks are trademarks of Facebook, Inc. TensorFlow, the TensorFlow logo and any related marks are trademarks of Google Inc. The OpenMP name and the OpenMP logo are registered trademarks of the OpenMP Architecture Review Board.

© 2021 Advanced Micro Devices, Inc. All rights reserved.

LEADING THE NEW ERA OF HPC



- Powered by AMD EPYC™ CPUs & AMD Instinct™ MI200 GPUs
- ~1.5 ExaFLOPS expected
- Expected to be more powerful than today's top 35 fastest supercomputers combined
- Shipment in 2021



- Powered by next gen AMD EPYC™ CPUs & next gen AMD Instinct™ GPUs
- ~2.0 ExaFLOPS expected
- Expected to be more powerful than today's 118 fastest supercomputers combined
- Shipment in 2023




LUMI

Nikhef



EuroHPC
Joint Undertaking





3rd Gen AMD EPYC™ PROCESSORS

New Features

THE WORLD'S HIGHEST PERFORMANCE
SERVER CPU*

“Zen3” Core Delivering Performance Enhancements

Enhanced Memory Performance

- Infinity Fabric™ and Memory Clock Synchronized
- Largest Available x86 L3 Cache; With Up To 32MB / Core

Enhanced Security Features

Supports 4, 6 or 8 Memory Channels Configs

Drop-in Compatible with EPYC 7002 Series Platforms

(BIOS update required)


OUR BET ON INNOVATIONS TO BEND TIME

ACCELERATED PLATFORM THAT IS UNIFIED, OPEN, AND OPTIMIZED FOR HPC



AMD
CDNA

DOMAIN-OPTIMIZED
ARCHITECTURE

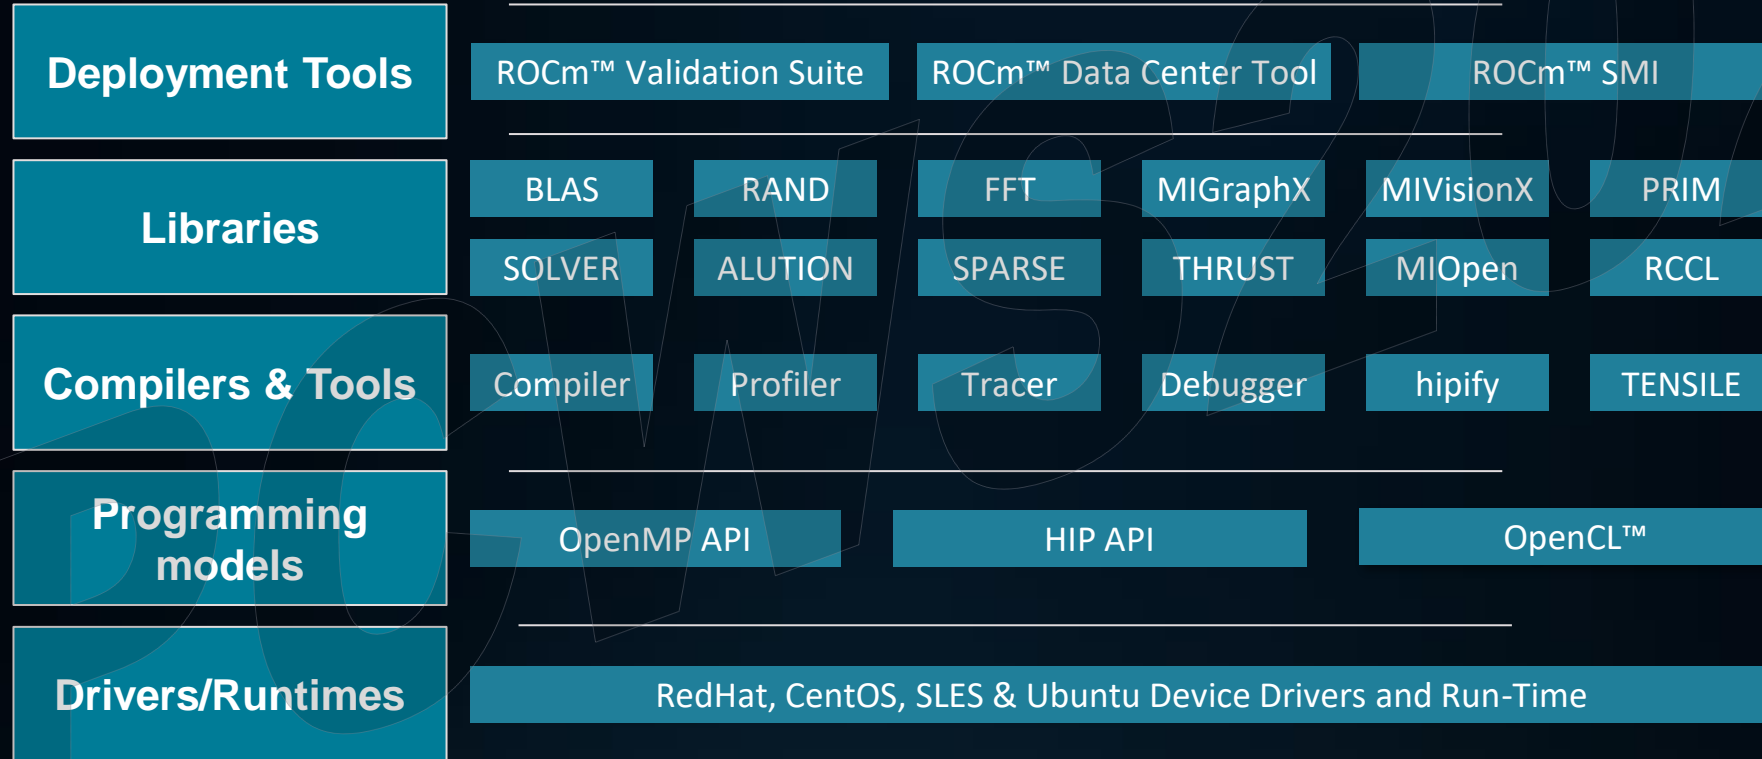


OPEN & PORTABLE
SOFTWARE



UNIFIED CPU & GPU
PLATFORM

AMD ROCm™ Platform



HIP – A common Kernel Language for Devices

CUDA VECTOR SAXPY FOR GPU

```
__global__
void saxpy(size_t n, float a, float * x, float * y)
{
    size_t i = threadIdx.x + blockIdx.x * blockDim.x;
    y[i] = a * x[i] + y[i];
}
```

```
// allocate the device memory
cudaMalloc(&x_dev, sizeof(*x_d) * count);
cudaMalloc(&y_dev, sizeof(*y_d) * count);

// copy stuff into the device buffers
cudaMemcpyAsync(x_d, x_h, sizeof(*x) * count,
                cudaMemcpyHostToDevice, stream);
cudaMemcpyAsync(y_d, y_h, sizeof(*y) * count,
                cudaMemcpyHostToDevice, stream);

saxpy<<<gridBlocks,blockThreads,0,stream>>>(n, a, x_d, y_d);
```

HIP VECTOR SAXPY FOR GPU OR CPU

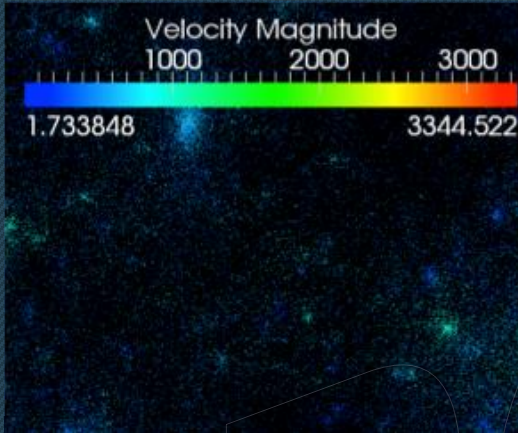
```
__global__
void saxpy(size_t n, float a, float * x, float * y)
{
    size_t i = threadIdx.x + blockIdx.x * blockDim.x;
    y[i] = a * x[i] + y[i];
}
```

```
// allocate the device memory
hipMalloc(&x_dev, sizeof(*x_d) * count);
hipMalloc(&y_dev, sizeof(*y_d) * count);

// copy stuff into the device buffers
hipMemcpyAsync(x_d, x_h, sizeof(*x) * count,
               hipMemcpyHostToDevice, stream);
hipMemcpyAsync(y_d, y_h, sizeof(*y) * count,
               hipMemcpyHostToDevice, stream);

saxpy<<<gridBlocks,blockThreads,0,stream>>>(n, a, x_d, y_d);
```

AMD ROCm™ ECOSYSTEM EXPANDING RAPIDLY WITH EASE OF PORTABILITY

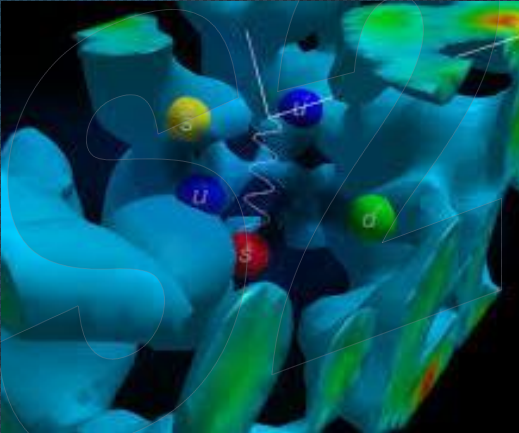


Velocity Magnitude
1000 2000 3000
1.733848 3344.522

Ported in an Afternoon
HACC
Cosmology



Ported in a Single Day
SPECFEM3D
Seismology



Ported & tuned in 21 Days
QUDA
Quantum Physics



Ported in a Couple of Days
CHOLLA
Astrophysics

NAMD
Scalable Molecular Dynamics

LAMMPS

kokkos

Nekbone

GROMACS
FAST. FLEXIBLE. FREE.

MILC

Chroma

TensorFlow

PYTORCH

GridTools

ALTAIR

SIRIUS


AMBER

PICongPU

CP2K

LSMS

AMD journey to exascale



AMD EPYC™ drives a wide spectrum of memory and compute intensive applications
AMD Instinct™ + AMD Infinity Fabric™ for a unified CPU & GPU platform
Solving computational intense challenges needs efficient hardware & software

[mathieu.gontier@AMD.com](mailto:mathieu.gontier@amd.com)

AMD 