



Arm as a Viable Architecture for HPC and AI

EPCC Workshop on
Efficient Computing for High Energy Physics

Dr. Nathan John Sircombe
nathan.sircombe@arm.com

18th February 2020

70%

of the world's population
uses Arm technology




arm

Not just mobile phones!

History of Arm in HPC

A Busy Decade



2011 Calxeda

- 32-bit ARmv7-A – Cortex A9



2011-2015 Mont-Blanc 1

- 32-bit Armv7-A
- Cortex A15
- First Arm HPC system




2014 AMD Opteron A1100

- 64-bit Armv8-A
- Cortex A57
- 4-8 Cores



2015 Cavium ThunderX

- 64-bit Armv8-A
- 48 Cores



2017 (Cavium) Marvell ThunderX 2



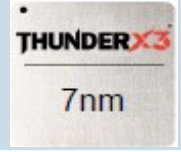














- 64-bit Armv8-A
- 32 Cores



2019 Fujitsu A64FX

- First Arm chip with SVE vectorisation
- 48 Cores

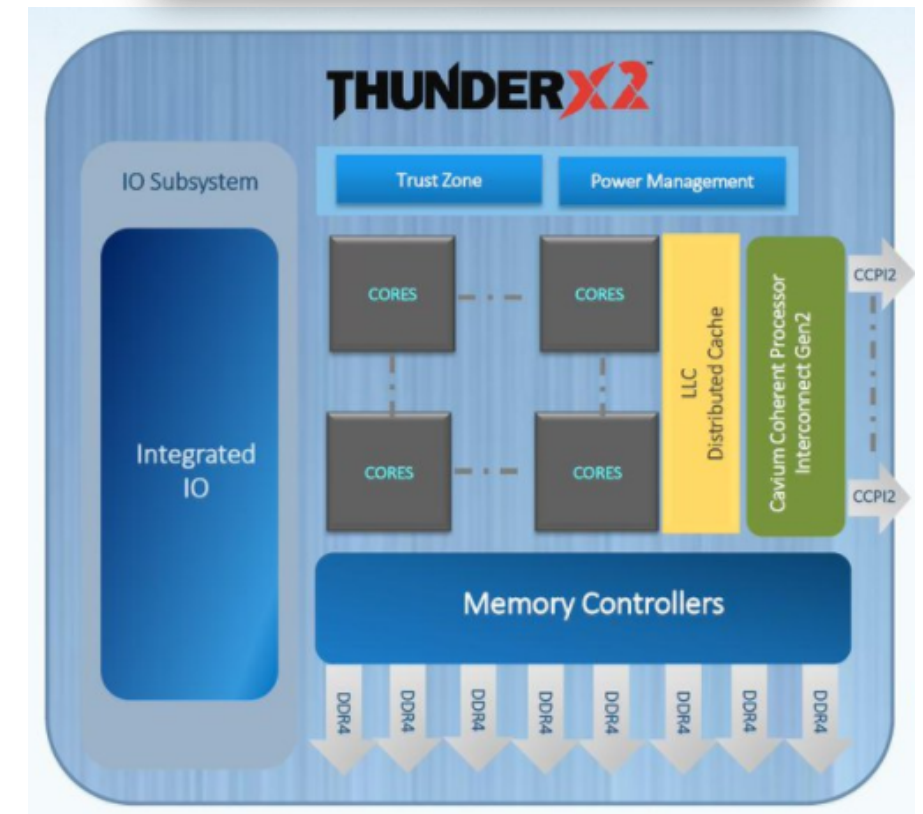
Variation in the Processor Market

Marvell (Cavium)							
Ampere (X-Gene)							
Fujitsu							
Huawei (HiSilicon)							
Amazon (Annapurna)							
EPI / SiPearl							
Other							

Marvell ThunderX2 CN99XX

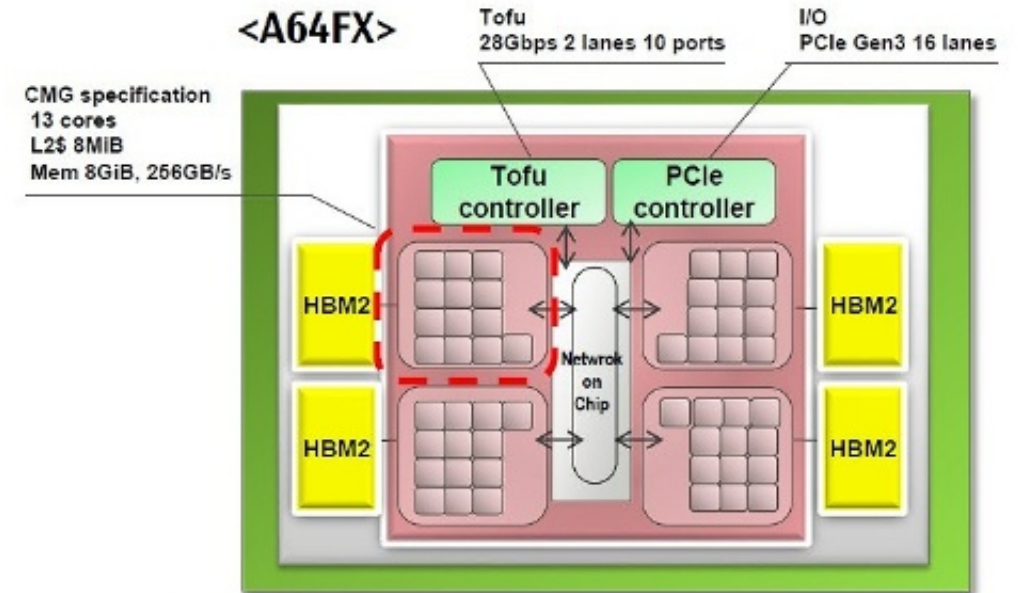


- Marvell's next generation 64-bit Arm processor
 - Taken from Broadcom Vulcan
- 32 cores @ 2.2 GHz (other SKUs available)
 - 4 Way SMT (up to 256 threads / node)
 - Fully out of order execution
 - **8 DDR4 Memory channels (~250 GB/s Dual socket)**
 - Vs 6 on Skylake
- Available in dual SoC configurations
 - CCPI2 interconnect
 - **180-200w / socket**
- Vector unit: **128-bit** NEON



Fujitsu A64FX

- Chip designed for RIKEN Fugaku (POST-K)
 - Based on Arm ISA technology
- 48 core 64-bit Armv8 processor
 - **+ 4 dedicated OS cores**
- **With SVE vectorisation**
 - **512 bit** vector length
- High performance
 - >2.7 TFLOPs
 - Low power : 15GF/W (dgemm)
- 32 GB HBM2
 - No DDR
 - **1 TB/s bandwidth**
- TOFU 3 interconnect



	A64FX (Post-K)	SPARC64 Xlfx (PRIMEHPC FX100)
ISA (Base)	Armv8.2-A	SPARC-V9
ISA (Extension)	SVE	HPC-ACE2
Process Node	7nm	20nm
Peak Performance	>2.7TFLOPS	1.1TFLOPS
SIMD	512-bit	256-bit
# of Cores	48+4	32+2
Memory	HBM2	HMC
Memory Peak B/W	1024GB/s	240GB/s x2 (in/out)

Deployments

- More Arm based CPUs are being adopted
 - Lots of large-scale deployments
- Different OEMs
 - Cray, HPE, Atos-Bull, Fujitsu, Huawei, E4
- EU Deployments
 - Isambard: Cray 10k TX2 cores
 - Catalyst 3 systems: HPE 4k TX2 core
 - **Future** Isambard 2: Cray A64FX
 - **Future** Deucalion: Cray A64FX



>5k ThunderX2 CPUs



2k Kunpeng 920 CPUs + 8k AI accelerators



150k+ Fujitsu A64FX CPUs



Deployments

Catalyst



Fulham Catalyst system at EPCC

- Deployments to accelerate the growth of the Arm **HPC** ecosystem
- Each machine has:
 - 64 HPE Apollo 70 nodes
 - Dual 32-core Marvell ThunderX2 nodes
 - 4096 cores per system
 - 256GB of memory / node
 - Mellanox InfiniBand interconnects
- OS: SUSE Linux Enterprise Server for HPC
- Signup for access:
 - <https://safe.epcc.ed.ac.uk/safadmin/>
 - Email olly.perks@arm.com for more information



Bristol: VASP, CASTEP, Gromacs, CP2K, Unified Model, NAMD, Oasis, NEMO, OpenIFS, CASINO, LAMMPS



EPCC: WRF, OpenFOAM, Two PhD candidates



Leicester: Data-intensive apps, genomics, MOAB Torque, DiRAC collaboration



A64FX Now in Top500 - #159

<https://www.top500.org/system/179706>

System	Year	Vendor	Cores	Rmax (GFlop/s)	Rpeak (GFlop/s)
A64FX prototype - Fujitsu A64FX, Fujitsu A64FX 48C 2GHz, Tofu interconnect D	2019		36,864	1,999,500	2,359,296

Green500 - #1

TOP500						Power
Rank	Rank	System	Cores	Rmax (TFlop/s)	Power (kW)	Efficiency (GFlops/watts)
1	159	A64FX prototype - Fujitsu A64FX, Fujitsu A64FX 48C 2GHz, Tofu interconnect D , Fujitsu Numazu Plant Japan	36,864	1,999.5	118	16.876

153W/node

- Prototype of Fugaku system
 - Fraction of the size of the final deployment
- Using Fujitsu software stack
 - Compiler and MPI



The Cloud

Open access to server class Arm

VERNE GLOBAL Verne Global: What is hpcDIRECT? In partnership with Arm and AWS

VERNE GLOBAL

in partnership with

arm

POWERED BY AWS GRAVITON2 PROCESSORS

M6g, R6g, C6g instances for EC2

New generation of Arm-based instances powered by AWS Graviton2 processors offer 40% better price/performance than current x86-based instances

M6g	R6g	C6g
PREVIEW TODAY	COMING SOON	COMING SOON



AWS Graviton Processor

First Arm Cloud Instances



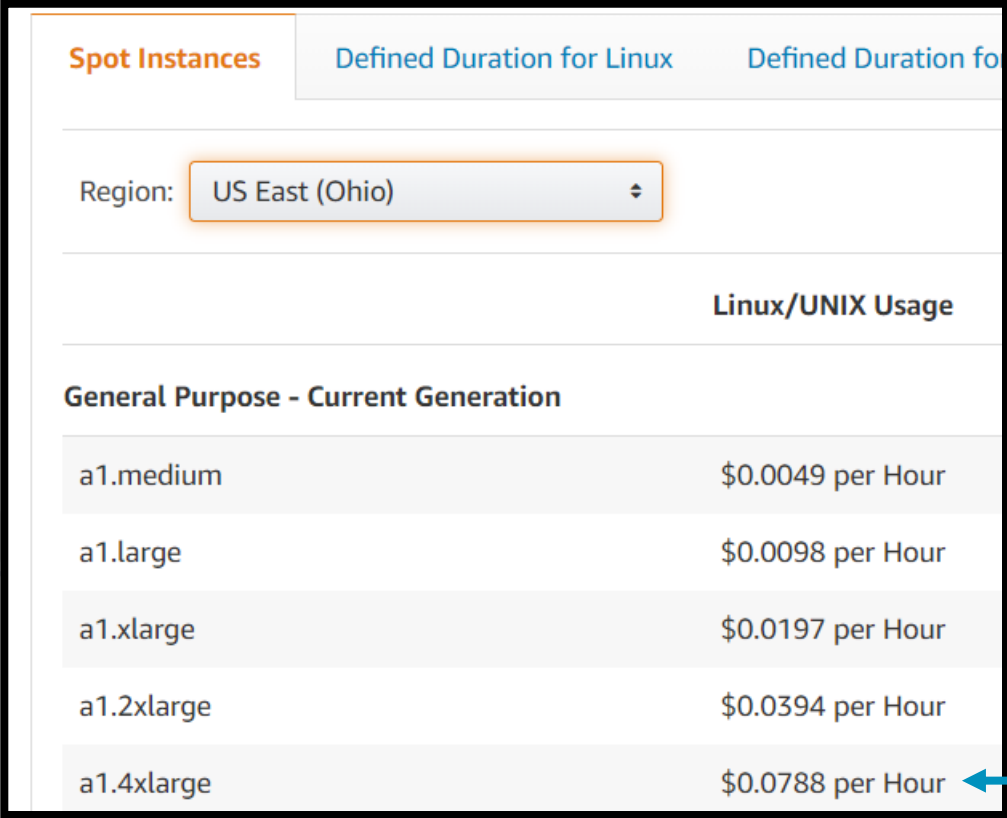
c1.large.arm

With 96 physical Arm cores, this server is anything but a lightweight - and it comes with 128 GB of RAM for just \$0.50/hr. Nice!



The Cloud

Open access to server class Arm



Spot Instances	
Defined Duration for Linux	
Defined Duration for Linux	
Region:	US East (Ohio)
Linux/UNIX Usage	
General Purpose - Current Generation	
a1.medium	\$0.0049 per Hour
a1.large	\$0.0098 per Hour
a1.xlarge	\$0.0197 per Hour
a1.2xlarge	\$0.0394 per Hour
a1.4xlarge	\$0.0788 per Hour

16 Cores A72 @ 2.3 GHz

With 96 physical Arm cores, this server is anything but a lightweight - and it comes with 128 GB of RAM for just \$0.50/hr. Nice!

arm

Software Ecosystem

Not Just Hardware

- Comprehensive software ecosystem
 - From Operating systems to Applications
 - Schedulers to file systems
- Everything you need to run an HPC service
- Vendor and OSS solutions



Functional Areas	Components include
Base OS	Centos, RHEL, Ubuntu, SUSE, SLES
Administrative Tools	Conman, Ganglia, Lmod, LosF, Nagios, pdsh, pdsh-mod-slurm, prun, EasyBuild, ClusterShell, mrsh, Genders, Shine, test-suite
Provisioning	Warewulf
Resource Mgmt.	SLURM, PBS Pro, Munge
I/O Services	Lustre client + server, NFS
Numerical/Scientific Libraries	Boost, GSL, FFTW, Metis, PETSc, Trilinos, Hypre, SuperLU, SuperLU_Dist, Mumps, OpenBLAS, Scalapack, SLEPc, PLASMA, ptScotch
I/O Libraries	HDF5 (pHDF5), NetCDF (including C++ and Fortran interfaces), Adios
Compiler Families	GNU (gcc, g++, gfortran), LLVM, Cray, Fujitsu, Arm
MPI Families	OpenMPI, MPICH, MVAPICH2, Cray, HPE
Development Tools	Autotools (autoconf, automake, libtool), Cmake, Valgrind, R, SciPy/NumPy, hwloc
Performance Tools	PAPI, IMB, pdtoolkit, TAU, Scalasca, Score-P, SIONLib

Cross-platform
debug and profile tools


Arm-only
Compiler and Libraries

arm ALLINEA STUDIO

Forge (DDT and MAP)
and Performance Reports
with support for Arm

Arm Fortran, C & C++ Compilers,
interoperable with Forge
and Performance Reports
Arm Performance Libraries

arm ALLINEA STUDIO




Fortran Compiler

- Fortran 2003 support
- Partial Fortran 2008 support
- OpenMP 3.1
- Directives to support explicit vectorization control
- SVE



C/C++ Compiler

- C++ 14 support
- OpenMP 4.5 without offloading
- SVE



Performance Libraries

- Optimized math libraries
- BLAS, LAPACK and FFT
- Threaded parallelism with OpenMP
- Optimized maths intrinsics



Forge

- Profile, Tune and Debug
- Scalable debugging with DDT
- Parallel Profiling with MAP



Performance Reports

- Analyze your application
- Memory, MPI, Threads, I/O, CPU metrics

Tuned by Arm for server-class Arm-based platforms

arm COMPILER

Commercial C/C++/Fortran compiler with best-in-class performance



Compilers tuned for Scientific Computing and HPC



Latest features and performance optimizations



Commercially supported by Arm

Tuned for Scientific Computing, HPC and Enterprise workloads

- Processor-specific optimizations for various server-class Arm-based platforms
- Optimal shared-memory parallelism using latest Arm-optimized OpenMP runtime

Linux user-space compiler with latest features

- C++ 14 and Fortran 2003 language support with OpenMP 4.5*
- Support for Armv8-A and SVE architecture extension
- Based on LLVM and Flang, leading open-source compiler projects

Commercially supported by Arm

- Available for a wide range of Arm-based platforms running leading Linux distributions – RedHat, SUSE and Ubuntu

arm PERFORMANCE LIBRARIES



Best-in-class performance



Commercially Supported
by Arm



Validated with
NAG test suite

- Commercial 64-bit ArmV8-A math Libraries
 - Commonly used low-level maths routines – BLAS, LAPACK and FFT
 - Optimised maths intrinsics
 - Validated with NAG's test suite, a de facto standard
- Best-in-class performance with commercial support
 - Tuned by Arm for specific cores - like TX2
 - Maintained and supported by Arm for wide range of Arm-based SoCs
- Silicon partners can provide tuned micro kernels for their SoCs
 - Partners can contribute directly through open source routes
 - Parallel tuning within our library increases overall application performance

Applications



Applications & frameworks

abinit, psdns, arbor, qmcpack, castep, quantumespresso, flecsale, raja, gromacs, sparta, kokkos, specfem3d, tensorflow, geant4, lammmps, sw4, pytorch, mxnet, nalu, milc, thornado, namd. vasp, nwchem, openfoam, wrf...



Mini apps

branson, pennant. cloverleaf, pf3dkernels, cloverleaf3d, quicksilver, e3smkernels, snap, kripke, snbone, lulesh.f tealeaf, miniamr, minife, minighost, nekbone. neutral...



Benchmarks

amg, nsimd, carmpl, nsimd-sve, clom, npb, elefun, polybench. epcc_c, stream, epcc_f, tsvc, umt, graph500, xsbench, hpcg, hpl hydrobench, ncar...

Community resources

<https://gitlab.com/arm-hpc/packages/wiki/>




Application Performance

Early Results from Astra


System has been online for around two weeks , incredible team working round the clock, already running full application ports and many of our key frameworks

Baseline: Trinity ASC Platform (Current Production (LANL/SNL)), dual-socket Haswell




Monte Carlo

1.60X




CFD Models

1.45X




Hydrodynamics

1.30X



Molecular Dynamics


1.42X




Linear Solvers

1.87X


EM (EMPIRE) Code on Astra




Single node performance results




UM scalability, up to 10,240 cores




NEMO scalability, up to 8,192 cores



OpenSBLI scalability, up to 10,240 cores




GROMACS scalability, up to 8,192 cores




Arm Performance Libraries – Leading BLAS performance

Arm Compiler for Linux 20.0 vs latest OpenBLAS vs latest BLIS



- High serial performance for BLAS level 3 routines, such as GEMMs also have class-leading parallel performance
- Shown is DGEMM on square matrices using 56 threads on a ThunderX2


ArmPL 20.0 FFT vs FFTW 3.3.8



Architecture Adoption: Community Engagement

Training Events and Hackathons

- Arm as a viable alternative to X86
- Needs to be easy to port to
 - Working codes and performant codes
- Team of field application engineers
 - Work with code teams
 - Educate, port and optimize
- Successful previous events
- Next event:
 - **Arm HPC User Group**
 - <https://a-hug.org/>
 - SVE Hackathon: **11th March**
 - Meeting: **12th-13th March**



Galaxy simulation in SWIFTsim computed on Arm Catalyst during DiRAC Hackathon

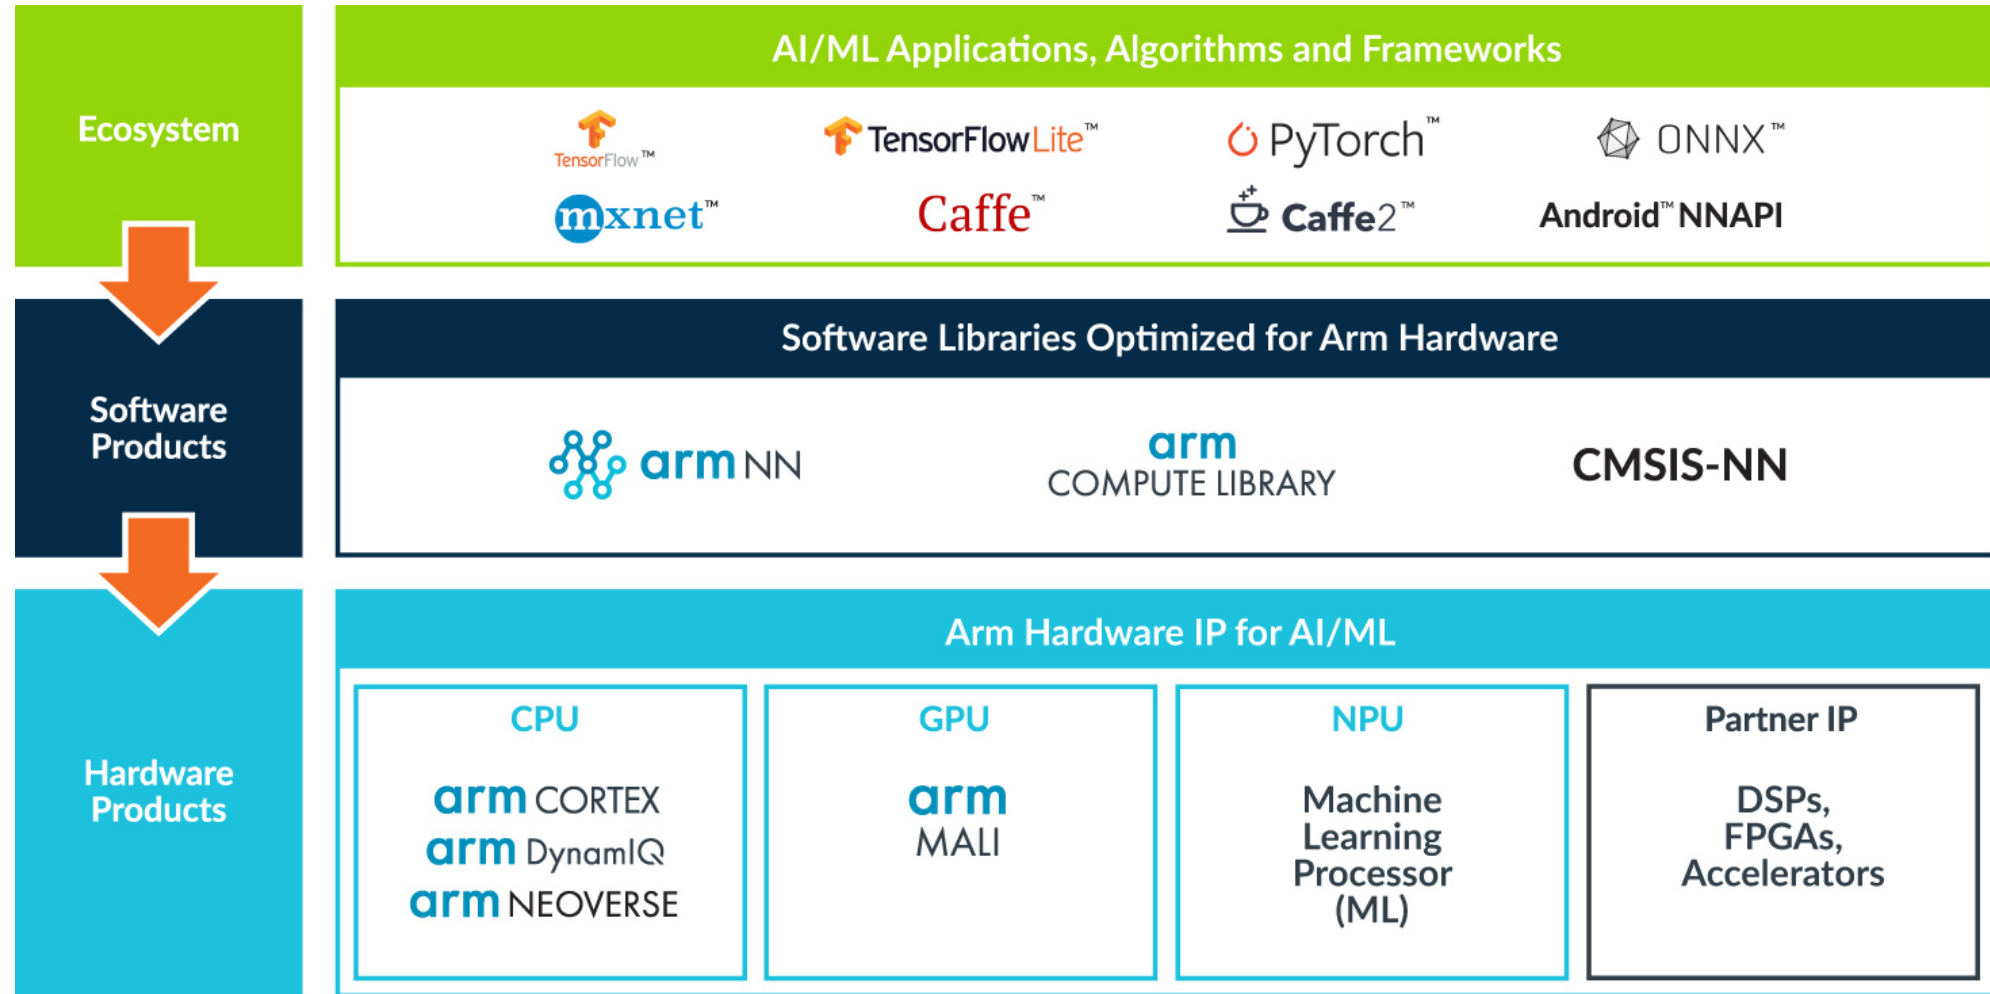


arm




Machine Learning and Artificial Intelligence

Machine Learning and Artificial Intelligence

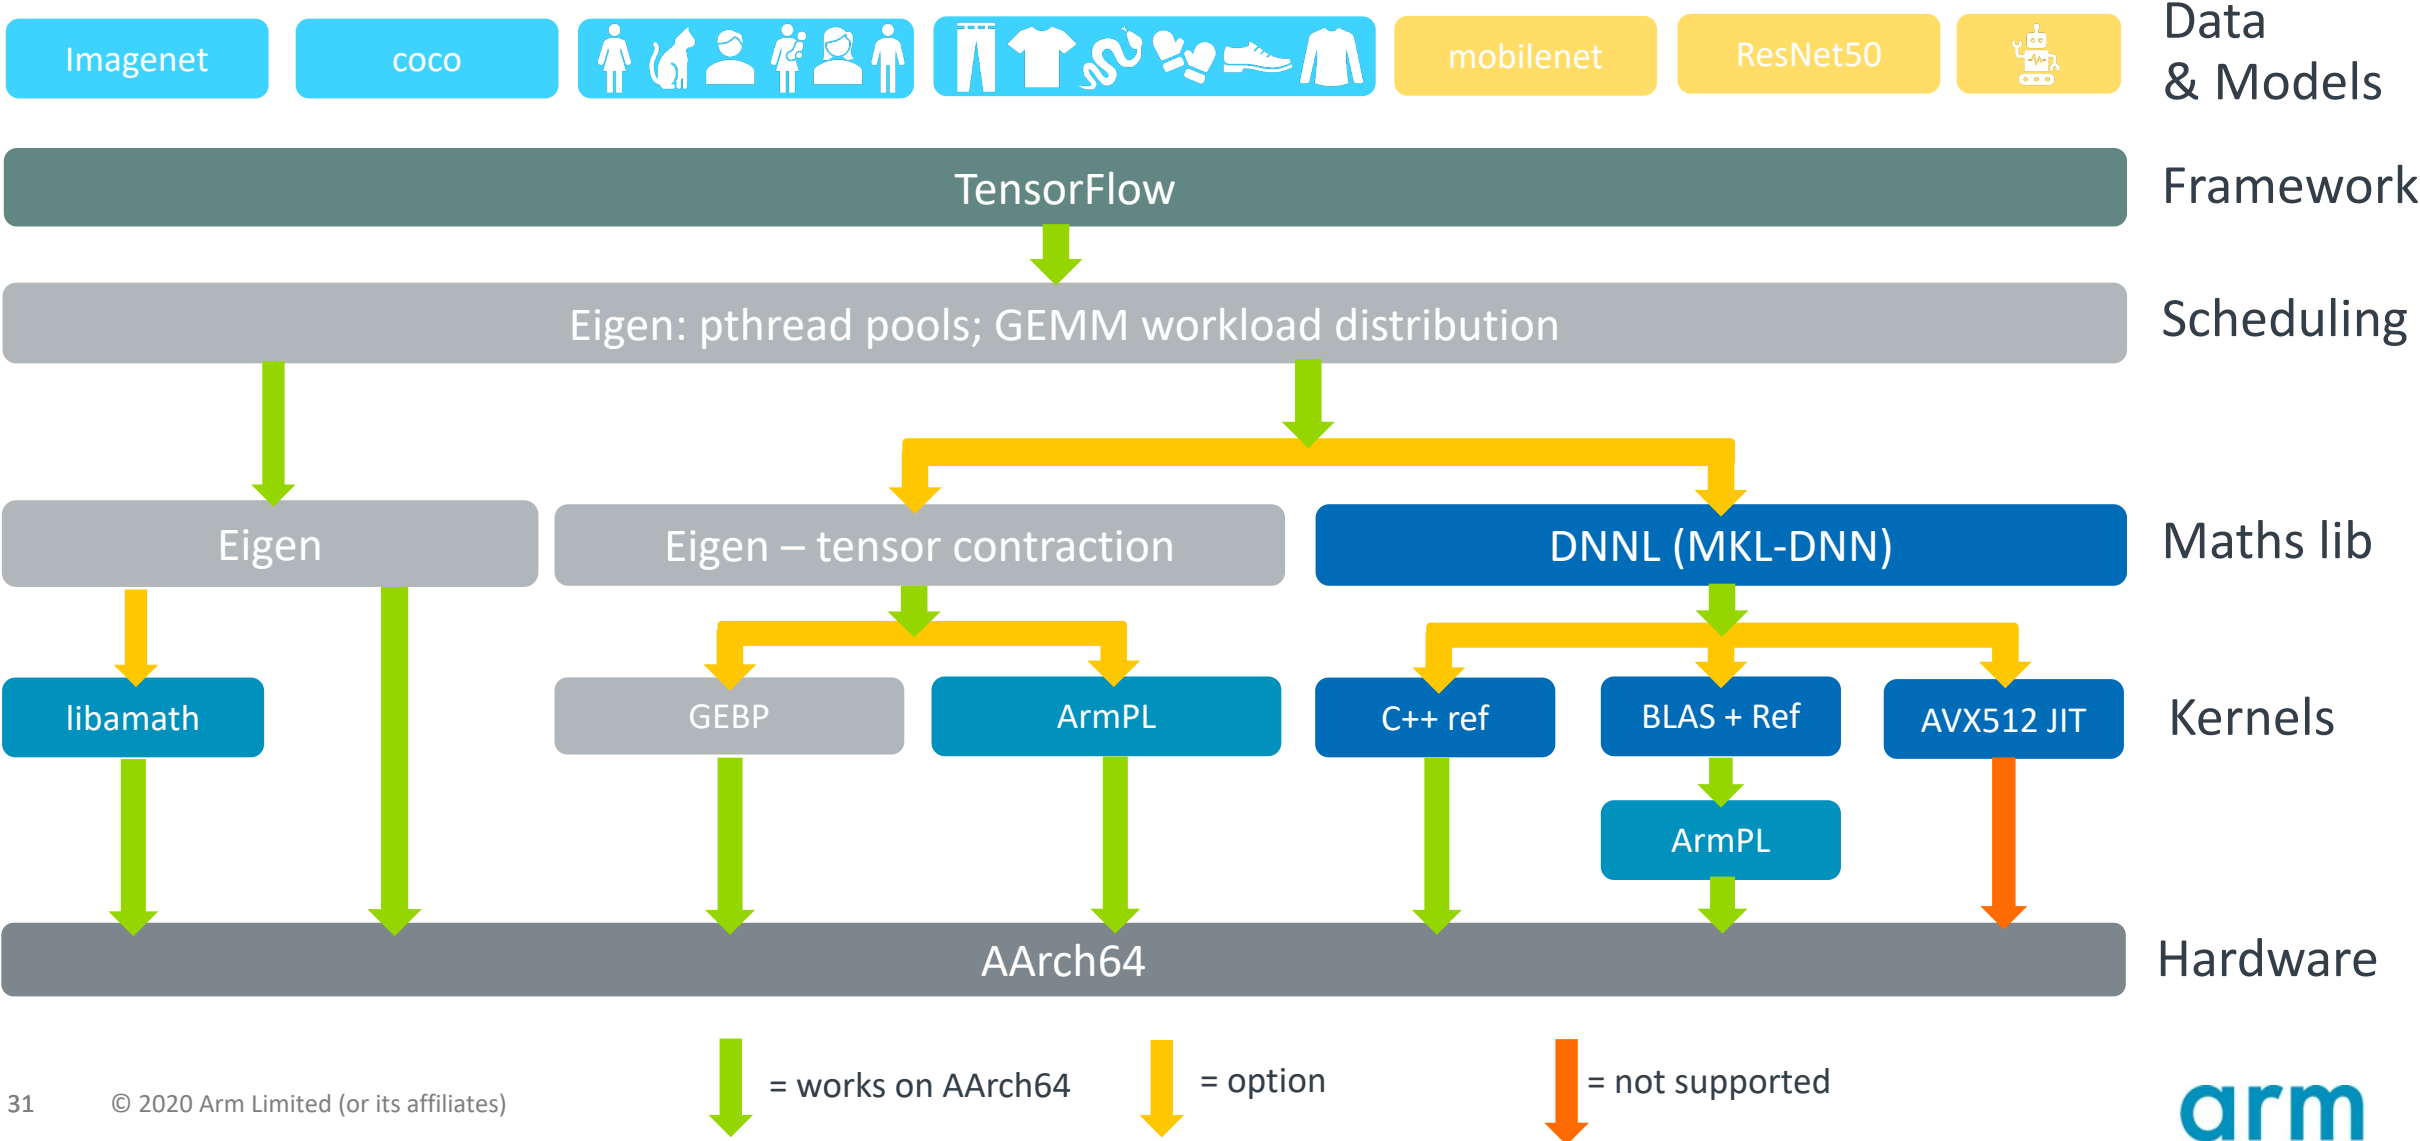


ML Frameworks on server-class Aarch64 platforms

- On-CPU server-scale ML workloads
- Leading frameworks and dependencies built on AArch64
 - TensorFlow: <https://gitlab.com/arm-hpc/packages/-/wikis/packages/tensorflow>
 - PyTorch: <https://gitlab.com/arm-hpc/packages/-/wikis/packages/pytorch>
 - MXNET: <https://gitlab.com/arm-hpc/packages/-/wikis/packages/mxnet>
- Docker tools for TensorFlow on GitHub
 - part of **ARM-software/Tool-Solutions**
 - <https://github.com/ARM-software/Tool-Solutions/tree/master/docker/tensorflow-aarch64>
 - Compiler: GCC 9.2
 - Maths libraries: Arm Optimized Routines and OpenBLAS 0.3.6
 - Python3 environment built from CPython 3.7 and containing:
 - NumPy 1.17.1
 - TensorFlow 1.15
 - TensorFlow Benchmarks



TensorFlow and maths libraries on AArch64




ML Frameworks on AArch64

- Focus has been on TensorFlow, and the maths libraries
 - Inference
 - Including ArmNN + ArmCL
 - Many-core systems
 - ArmCL note developed for > 8 cores, and doesn't scale well as-is
 - Significant GEMM, and vector maths, work
- Scope for improvements to performance and parallelism
- We're actively working on:
 - Optimized kernels
 - Improved scaling on many-core SoCs
 - Better support for AArch64 Neon and SVE
 - Leveraging Arm Performance Libraries for HPC-ML workloads
 - Enablement of AArch64 support in key libraries
 - Provision of OS implementations of key kernels

ML in HPC

- Catalyst cluster located in Leicester, 2 x Cavium ThunderX2(R) CPU CN9980 v2.1 @ 2.20GHz per node
- Distributed training benchmarks run at scale on Catalyst system
 - Cosmoflow
 - ResNet101
 - Climate Segmentation
 - <https://github.com/sparticlesteve/climate-seg-benchmark>
 - DeepLabv3+NN, training via Synchronous SGD
- Work supported by DiRAC Post-Doctoral industrial placement

Climate Segmentation scaling on Catalyst



arm

Going Forward

The Future of Arm in HPC

What's next?

Processors

- By more vendors
 - Marvell, Ampere, Amazon, HiSilicon, Fujitsu
- Targeting different market segments
- All built on the Arm ecosystem
- Supported by the tools

Deployments

- Large scale and small scale deployments
- Increased exposure to the architecture
- More applications and libraries ported to Arm
 - Including ISVs
- Growing community

Commitment from Arm

- Neoverse – IP roadmap for silicon vendors
- Investment in software ecosystem
 - E.g. Flang / F18
- Support for customers
 - Applications
 - Software
 - Performance

Arm HPC Ecosystem

Get involved

www.arm.com/hpc

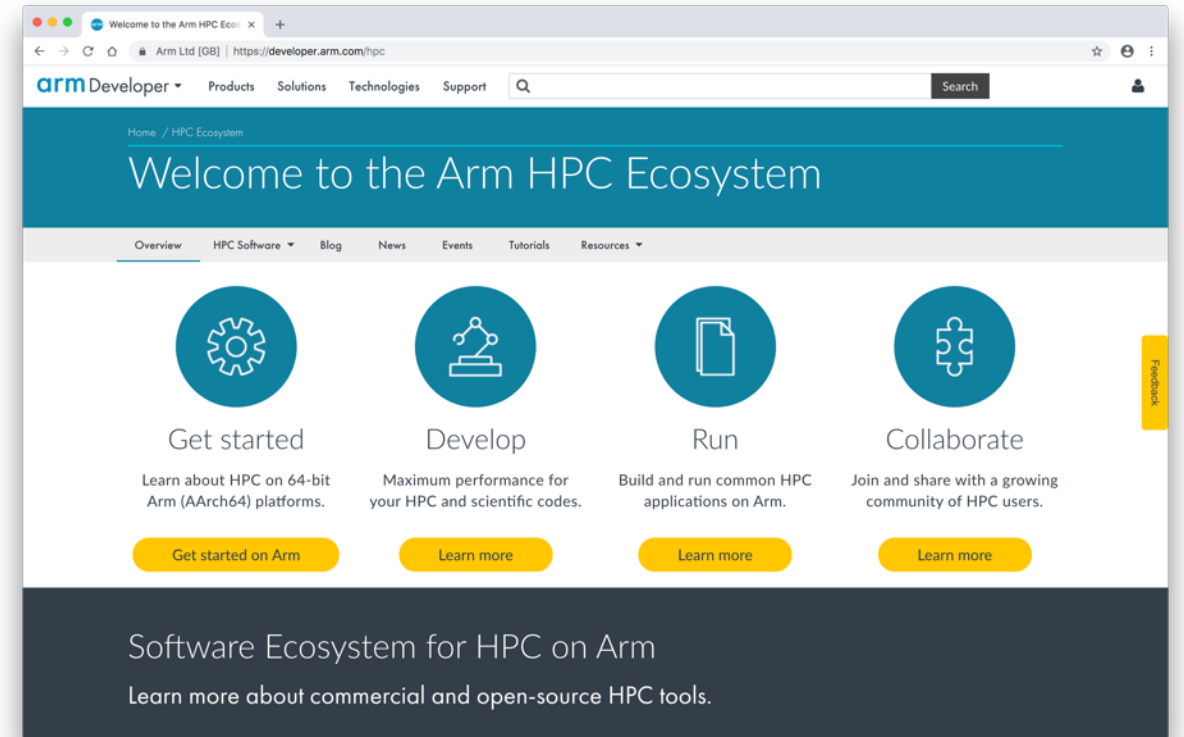
- News, events, blogs, webinars, etc.
- Quick-start guides for tools

www.gitlab.com/arm-hpc/packages/wikis

- Community collaboration site
- Guides for porting HPC applications

www.a-hug.org

- Arm HPC Users Group (AHUG)





The Arm trademarks featured in this presentation are registered trademarks or trademarks of Arm Limited (or its subsidiaries) in the US and/or elsewhere. All rights reserved. All other marks featured may be trademarks of their respective owners.

www.arm.com/company/policies/trademarks