LS-DYNA and High-Performance Computing

Jason Wang, Eric Day

Ansys

Outline

- *LS-DYNA x86_64 binaries*
- *CISC vs RISC*
- *ARM64 HPC performance and coming new CPUs*
- *LS-DYNA performance between releases*
- *Conclusion*

LS-DYNA x86_64 binaries

Most of the releases - ifort + MKL

- Runs on both Intel Xeon and AMD EPYC chips
- Produces identical results from same input on both chips for explicit

(MKL needs special environment variable for implicit)

Additional release - AOCC + AOCL

- Runs on both Intel Xeon and AMD EPYC chips
- Produces identical results from same input on both chips for explicit

(AOCL needs special environment variable for implicit)

Performance comparison among 4 different builds,

ifort (AVX2, AVX512) and aocc(AVX2, AVX512)

Ansys

Explicit - car2car (2.4 million elements, 30 ms)

LS-DYNA Binaries

binary	Compiler options
ifort 2019.6.324 avx2	-march=core-avx2 -mtune=core-avx2 -align array32byte
aocc 4.0.0 avx2	-mavx2
ifort 2019.6.324 avx512	-march=skylake-AVX512 -mtune=skylake-AVX512 -align array64byte -qopt-zmm-usage=high
aocc 4.0.0 avx512	-mavx512f

- Tested with Intel MPI
- No hardware dependent options but instruction set dependent options
- Produce same numerical results from different generations/brands of CPUs

AMD EPYC 9654 (Zen4/192 cores) : car2car

Number of ranks	Intel MPI runtime options
48	-genv I_MPI_PIN_PROCESSOR_LIST=allcores:shift=4 -np 48 -ppn 192
96	-genv I_MPI_PIN_PROCESSOR_LIST=allcores:shift=2 -np 96 -ppn 192
144	-genv I_MPI_PIN_PROCESSOR_LIST=allcores:shift -np 144 -ppn 192
192	-genv I_MPI_PIN_PROCESSOR_LIST=allcores:shift -np 192 -ppn 192

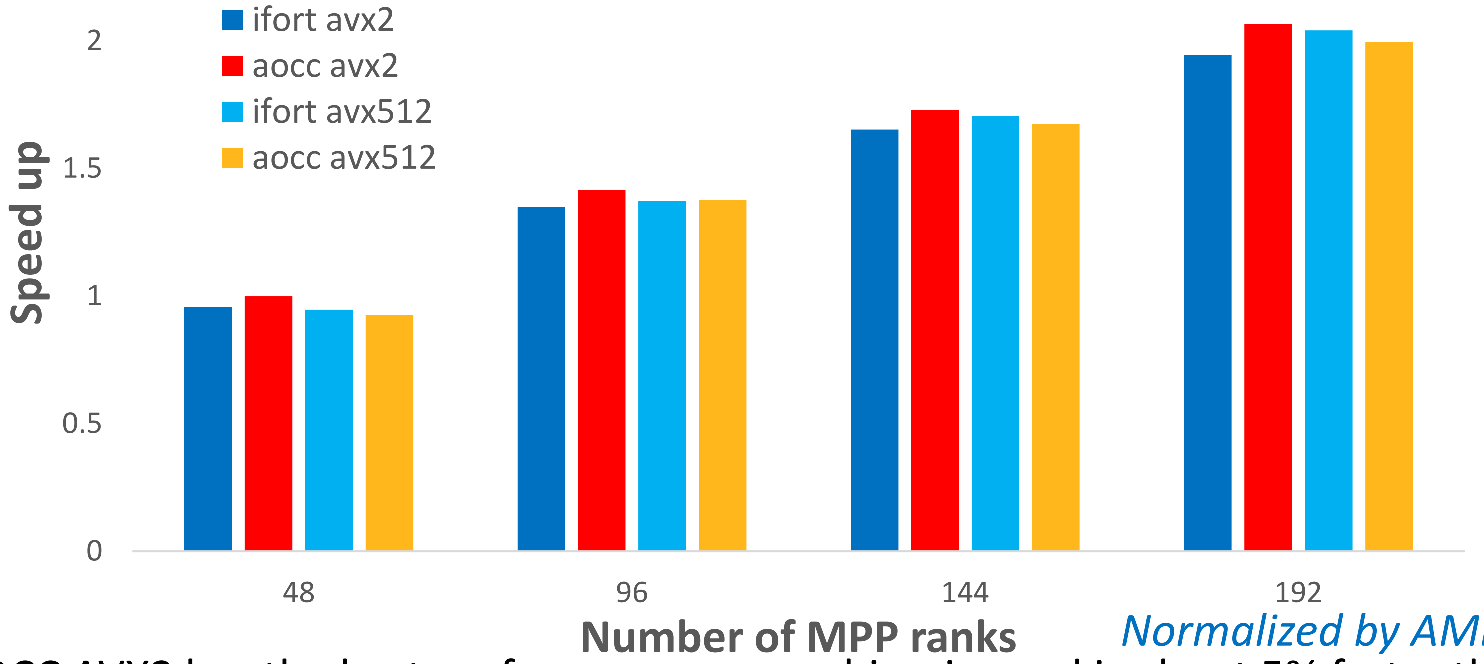

Normalized by AMD aocc avx2 48-core timing

- AOCC AVX2 has the best performance among binaries and is about 5% faster than ifort AVX2.
- Ifort AVX512 is a little faster than ifort AVX2
- Zen4 is the first AMD chip support AVX512.

Intel Xeon Gold 1642 (48 cores): car2car

Number of ranks	Intel MPI runtime options
24	-genv I_MPI_PIN_PROCESSOR_LIST=allcores:shift=2 -np 24 -ppn 48
48	-genv I_MPI_PIN_PROCESSOR_LIST=allcores:shift -np 48 -ppn 48

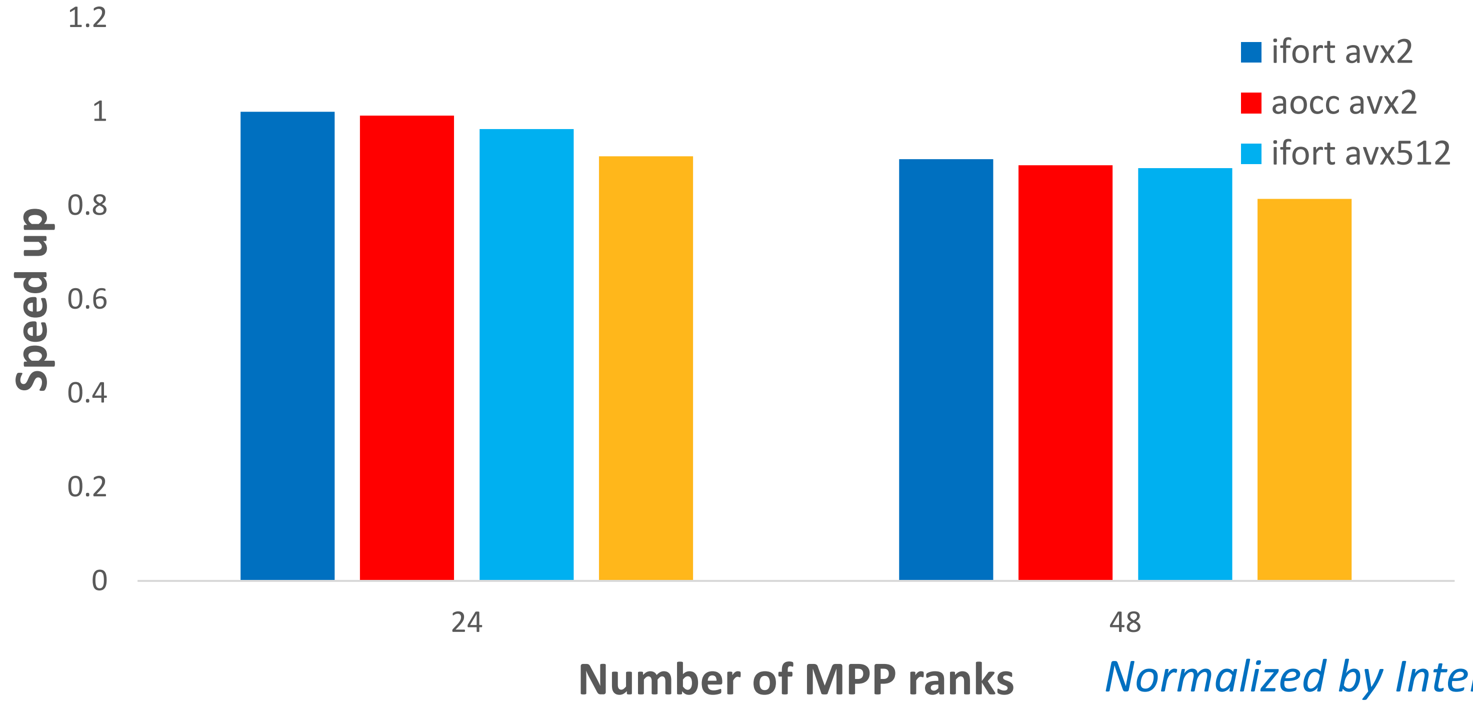

Normalized by Intel ifort avx2 24-core timing

- Ifort AVX2 has the best performance among 4 binaries and is about 1% faster than AOCC AVX2.
- AVX512 has less performance than AVX2 for both compilers
- AOCC AVX512 does not perform well on Intel chips.

Ansys

Numerical consistency between AMD Zen and Intel Xeon Scalable

Identical results for explicit analysis

- Same decomposition
- Set lstc_reduce on

- Additional setting for Intel MPI (2019 and above)

  -genv I_MPI_CBWR=2

  https://cdrdv2-public.intel.com/671217/mpi-dev-ref-lin-u6.pdf

Implicit – Cycl1e6 (1 million elements)

CPU

- AMD EPYC 9654
- Intel Xeon Gold 1642

Normalized by the best timing on each hardware

	AMD Zen -96 ranks		Intel Scalable - 32 ranks	
	Extra environment variable	**Speedup**	**Extra environment variable**	**Speedup**
AOCC/MKL	MKL_DEBUG_CPU_TYPE=5	0.97	MKL_DEBUG_CPU_TYPE=5	0.82
		0.44		0.82
AOCC/AOCL	BLIS_ARCH_TYPE=zen4	0.97	BLIS_ARCH_TYPE=zen4	0.83
		0.95		0.24
IFORT/MKL	MKL_DEBUG_CPU_TYPE=5	1	MKL_DEBUG_CPU_TYPE=5	0.86
		0.92		1

- Ifort/MKL has the best performance on both hardware ***(different environment variable setting)***
- Without proper environment variable, ***MKL and AOCL perform poorly***.
- Will release AOCC/AOCL and IFORT/MKL.***(AOCC/MKL for internal testing only)***
- AMD mentioned AOCL does not need the flag in the future release.

The Clash of Architectures: CISC vs. RISC

CISC - Complex Instruction Set Computer (**x86_64** Architecture)

- Intel Xeon, AMD EPYC,etc.
- Extensive instruction set
- Substantial computing power

RISC - Reduced Instruction Set Computer (**ARM64** Architecture)

- ARM64, RISC-V, IBM power, MIPS, SPARC, PA, etc.
- Base-level simple instructions
- Require multiple instructions to complete complex tasks
- Reduced power consumption

X86_64 Dominance in Datacenters/Cloud

The Longstanding Presence of x86_64 in Datacenters			
Raw Power		Reliability	
Well-established Optimizations	Vast Scientific Software	Mature Compilers	The Conventional Choice

2021Q2 Server Shipments by CPU Type

“Potential users will look at this ARM CPU - see that it is not faster than Intel on a per thread basis and is not x86-64 compatible and will turn away with a shrug. A minor price difference for a complete server is not enough to justify the risks of going from x86-64 to ARM.”

Ansys

Assessed ARM Processors - Specifications

Processor	Year Released	Cores, Frequency	Architecture, Microarchitecture	SIMD	Tech	Memory	Memory BW
Cavium ThunderX2 CN9975	2018	28 @2.4GHz	ARMv8.1-a, Vulcan	128bit Neon	16nm	8xDDR4-2666	170GB/s
Fujitsu A64FX	2019	48 @2.0GHz	Armv8.2-a, A64FX	256bit SVE	7nm		1,024GB/s
Ampere Altra Q64-30	2021	64 @3.0GHz	ARMv8.2-a, Neoverse-N1	2x128bit Neon	7nm	8xDDR4-3200	204GB/s
Apple M1 Max	2021	8 @3.2GHz + 2 efficiency	ARM 8.5-a, Firestorm	128bit Neon	5nm	16xLP DDR5-6400	408GB/s
AWS Graviton2	2020	64 @2.5GHz	ARMv8.2-a, Neoverse-N1	2x128bit Neon	7nm	8xDDR4-3200	204GB/s
AWS Graviton3	2022	64 @2.6GHz	ARMv8.4-a, Neoverse-V1	2x128bit Neon and 2x256bit SVE	5nm	8xDDR5	300GB/s

Neon

- 500,000 elements
- Run up to 8 cores
- Run on a single socket and NUMA node

<u>LS-DYNA Development source</u>

- MPP single precision
- OpenMPI 4.x
- ARM64 – armflang22.0.2 and gcc
- x86_64 – avx2, ifort190 and gcc

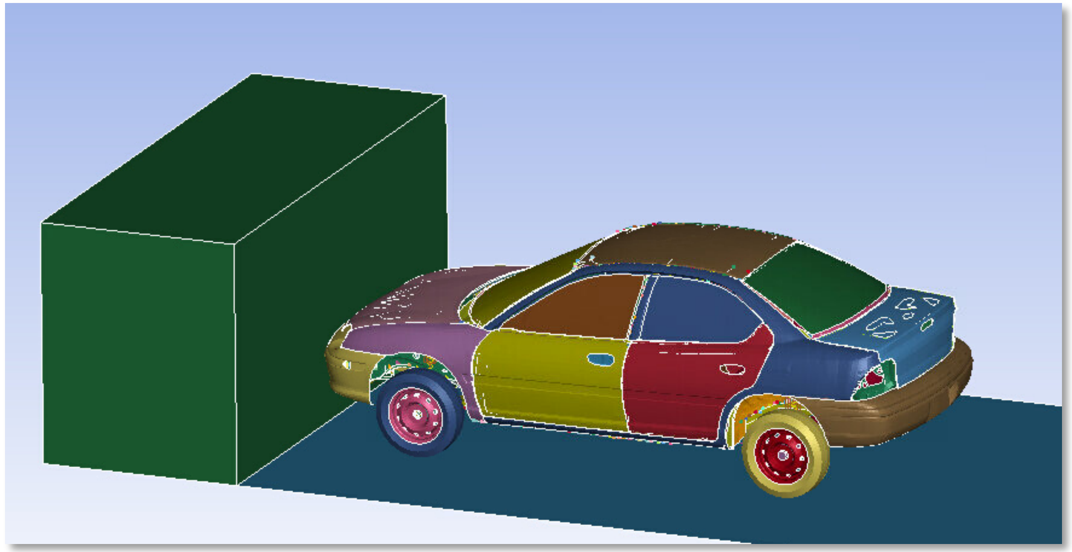

Neon – Performance by CORE

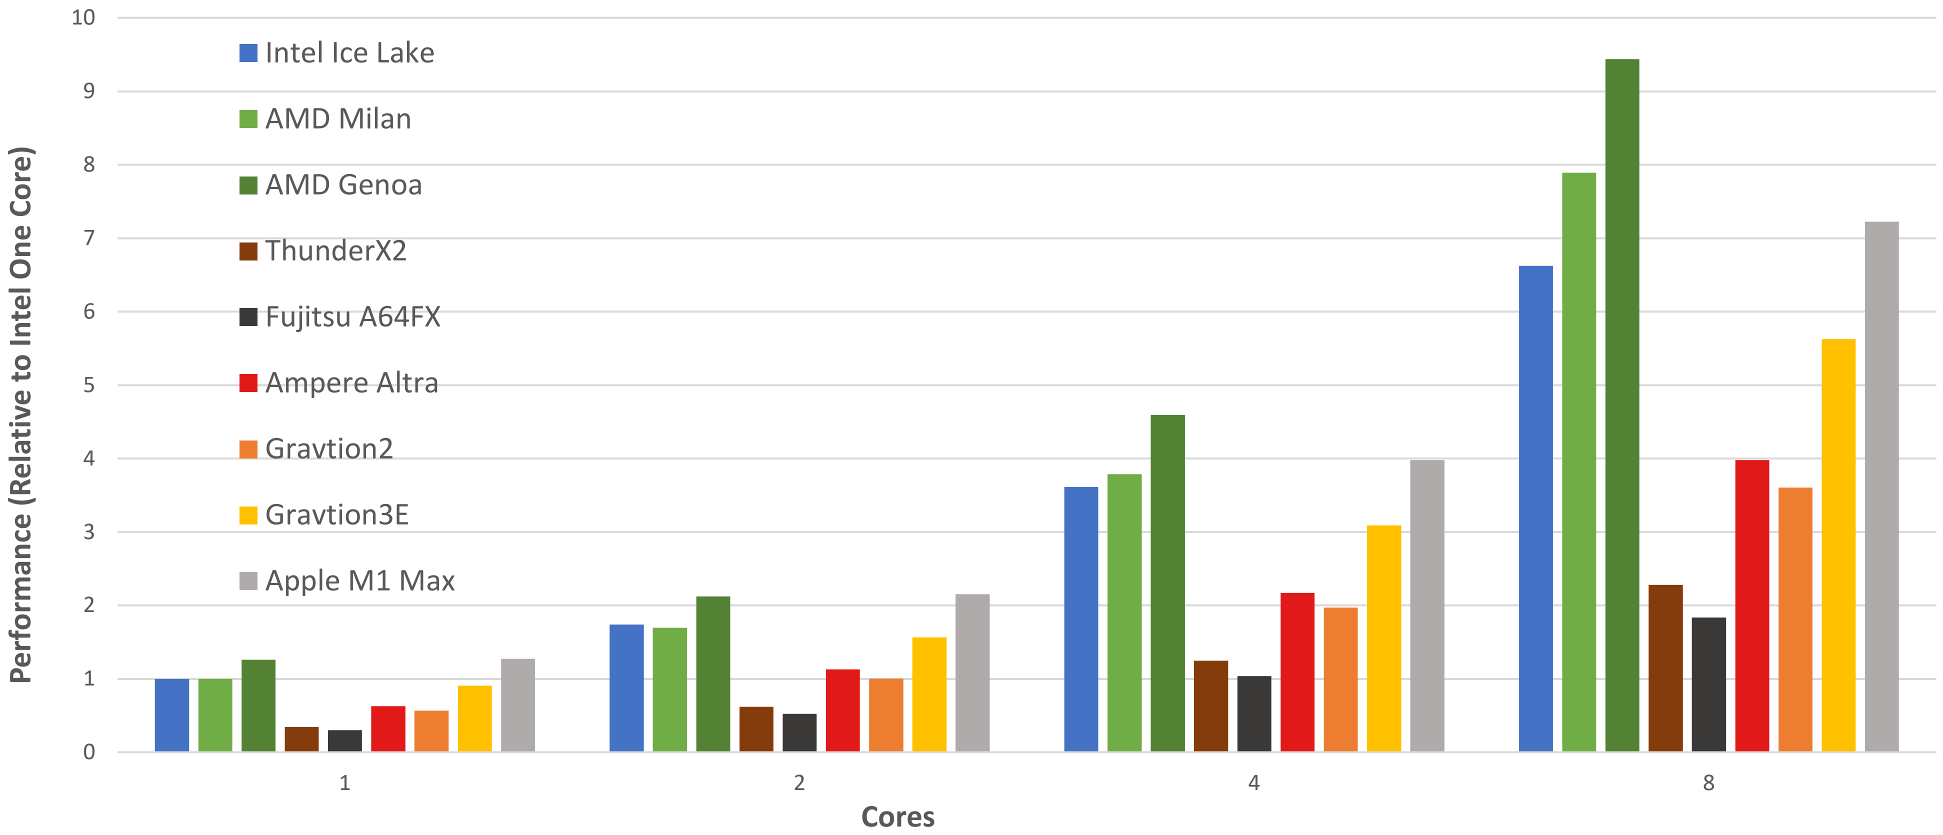

- Apple M1 max achieves exceptional single and double-core performance
- Contemporary ARM processors have good performance, but computational power may still lag behind the latest 4th generation EPYC and Xeon processors.

ODB-10M

- 10 million elements
- Run up to 8 nodes

LS-DYNA R12.1.0

- MPP single precision
- Multithreading disabled
- OpenMPI 4.x
- Amazon instances run with EFA

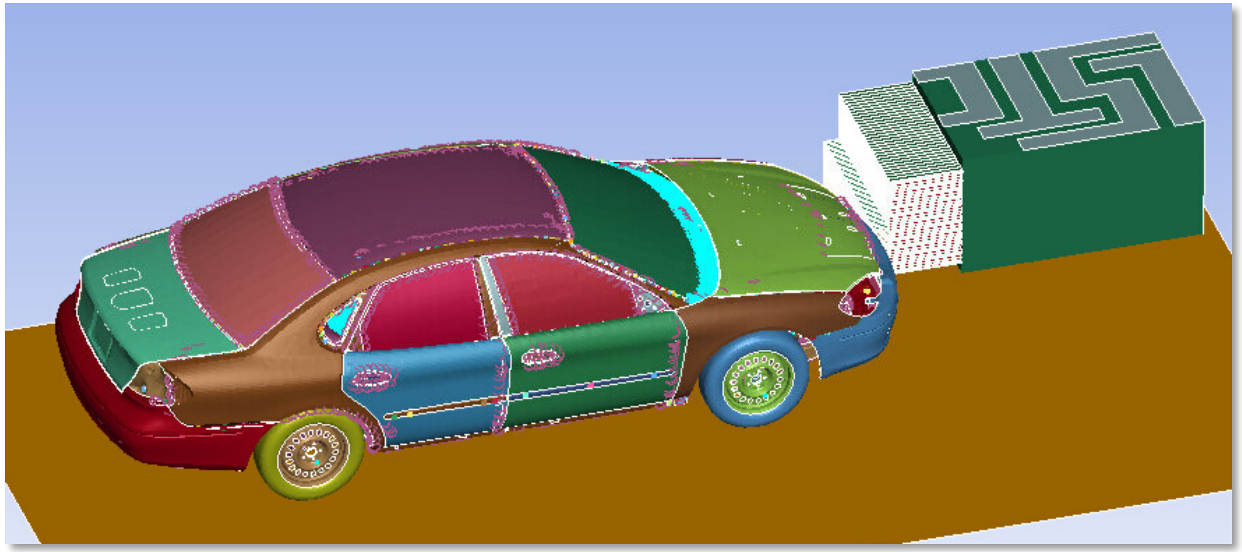

ODB-10M – Performance by NODE

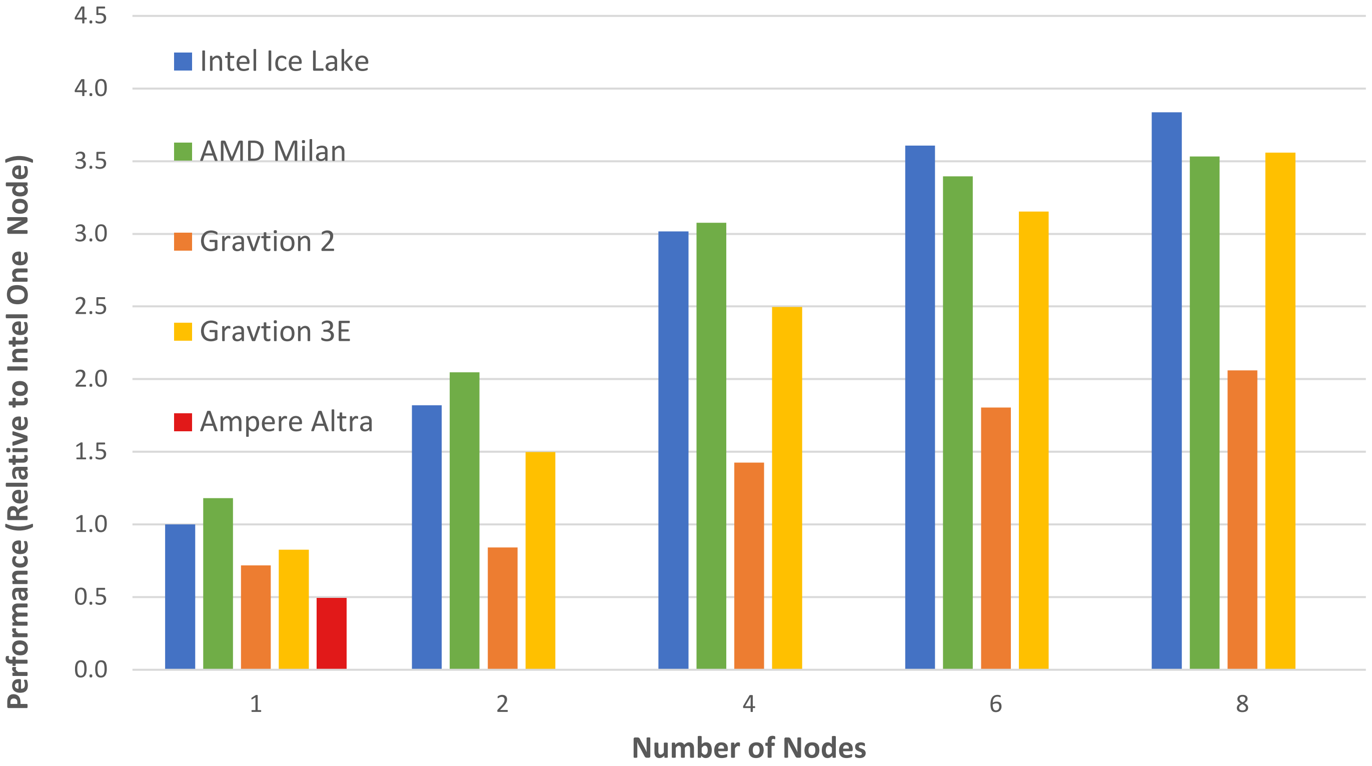

CPU	Cores/CPU	Sockets	Cores/node	memory channel/socket
Intel Ice Lake 8375C	32	2	64	8
AMD Milan 7R13	48	2	96	8
Ampere Altra Q64-30	80	1	80	8
AWS Gravtion2	64	1	64	8
AWS Graviton3E	64	1	64	8

- Most of automotive users use less than 384 cores/job (2-4 nodes).
- X86-64 has dual sockets – memory bandwidth/core is higher than arm64

ODB-10M – Price Performance

Environmental and Economical influence

A high price-performance instance offers the greatest computational capacity for every dollar invested

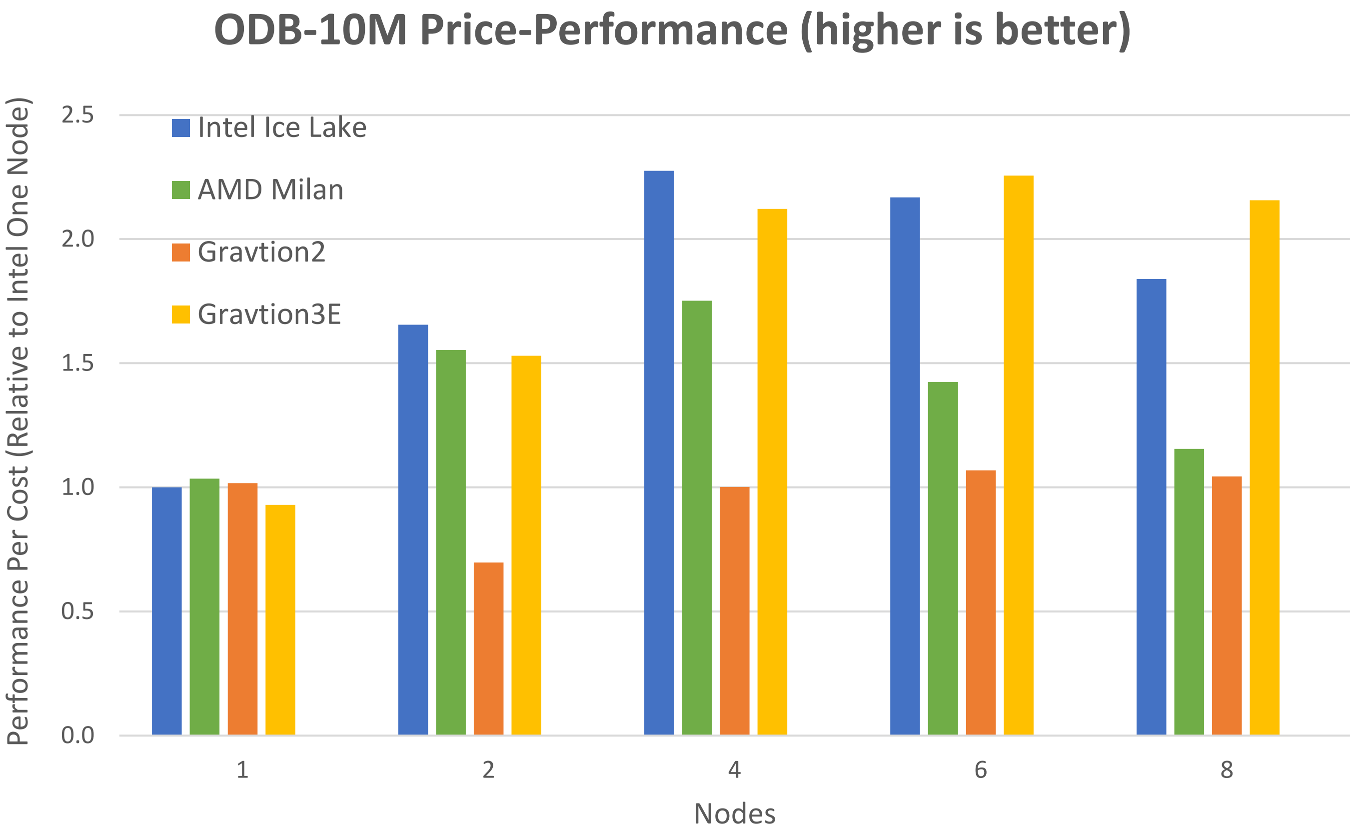

CPU	Cores/CPU	Sockets	TDP per CPU
Intel Ice Lake 8375C	32	2	300W
AMD Milan 7R13	48	2	225W
Ampere Altra Q64-30	80	1	180W
AWS Gravtion2	64	1	est. 110 - 130W
AWS Graviton3E	64	1	est. 210 - 295W

ARM **V1-platform** (Graviton3) prioritizes **computational power** over area and energy efficiency.

Future of ARM

Imminent Processors

- Ampere One – ARMv8.6, 192 cores. 4.3x VM's per rack compared to Intel 4th gen
- Nvidia Grace – ARMv9, 144 Neoverse-V2 cores.
- Nvidia Grace Hopper – CPU + GPU coherent memory

Further Innovations

- TSMC 3nm platform – rumored for Apple M3, Graviton4
- OpenMPI 5.0
- Expanding software libraries for arm64

ARM's 2028 Server Market Share Target

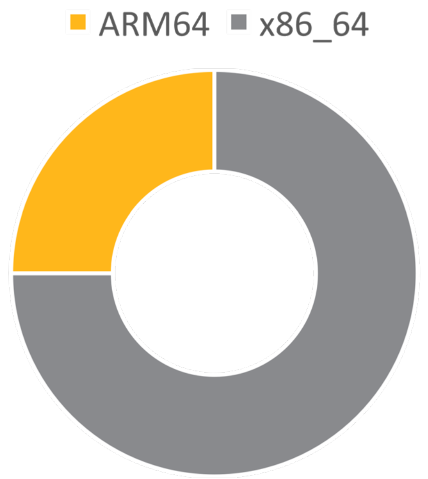

Performance between releases

- Model: ODB-10M
- AMD EPYC 7V73X 64-Core Processor (2 sockets, 120 total cores)
- Using AVX2 binaries
- Performance Relative to R11.2.2 single node (higher is better)

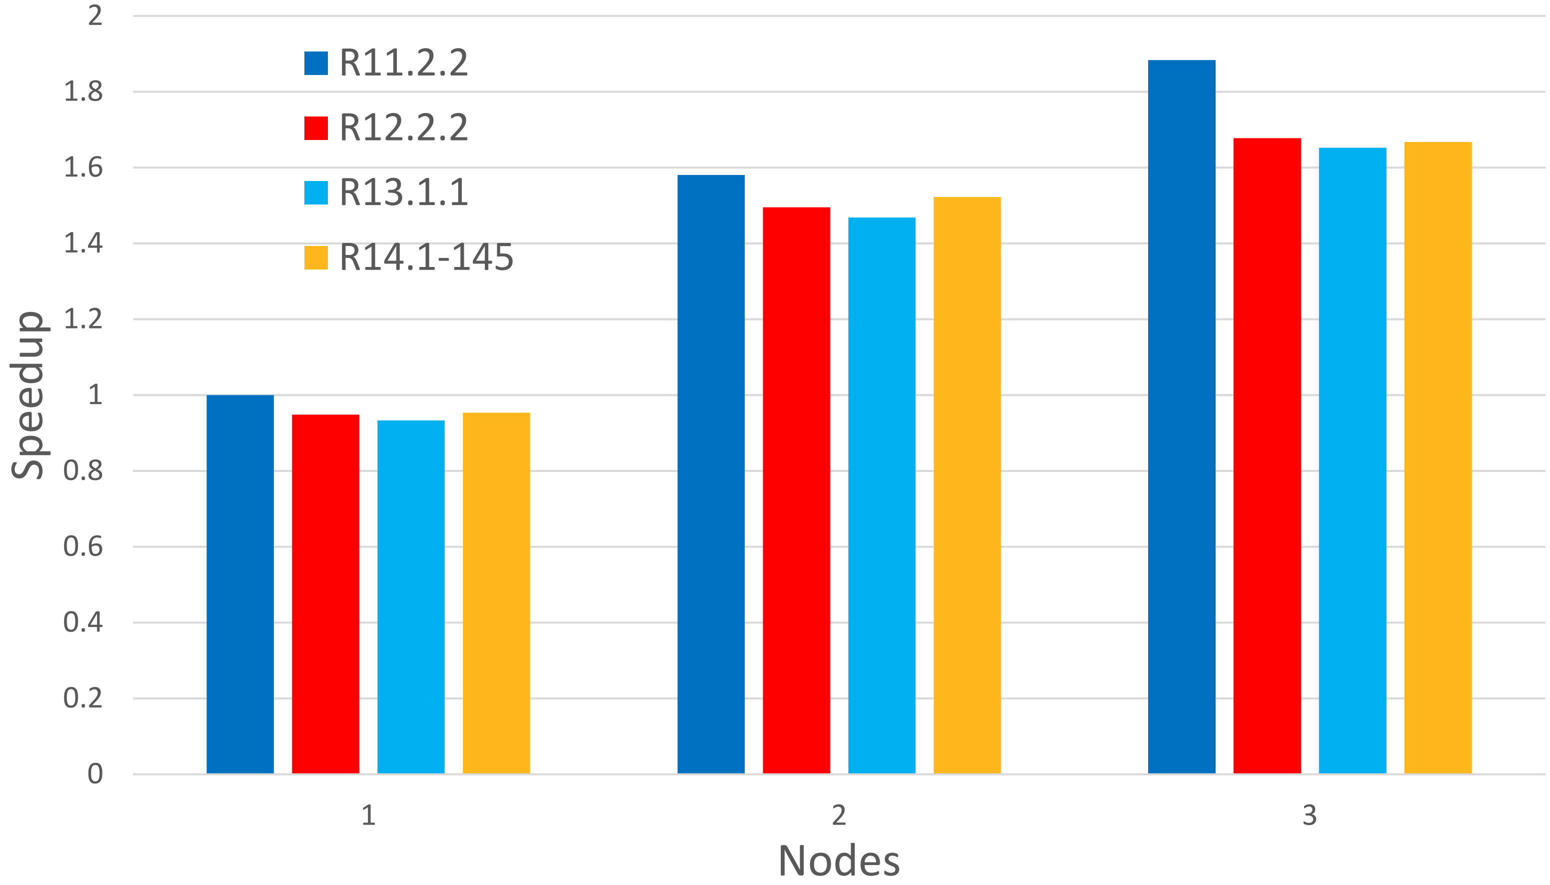

Performance between releases

- Model: ODB-10M
- Intel Xeon Platinum 8168 CPU (2 sockets, 44 total cores)
- Using AVX2 binaries
- Performance Relative to R11.2.2 single node (higher is better)

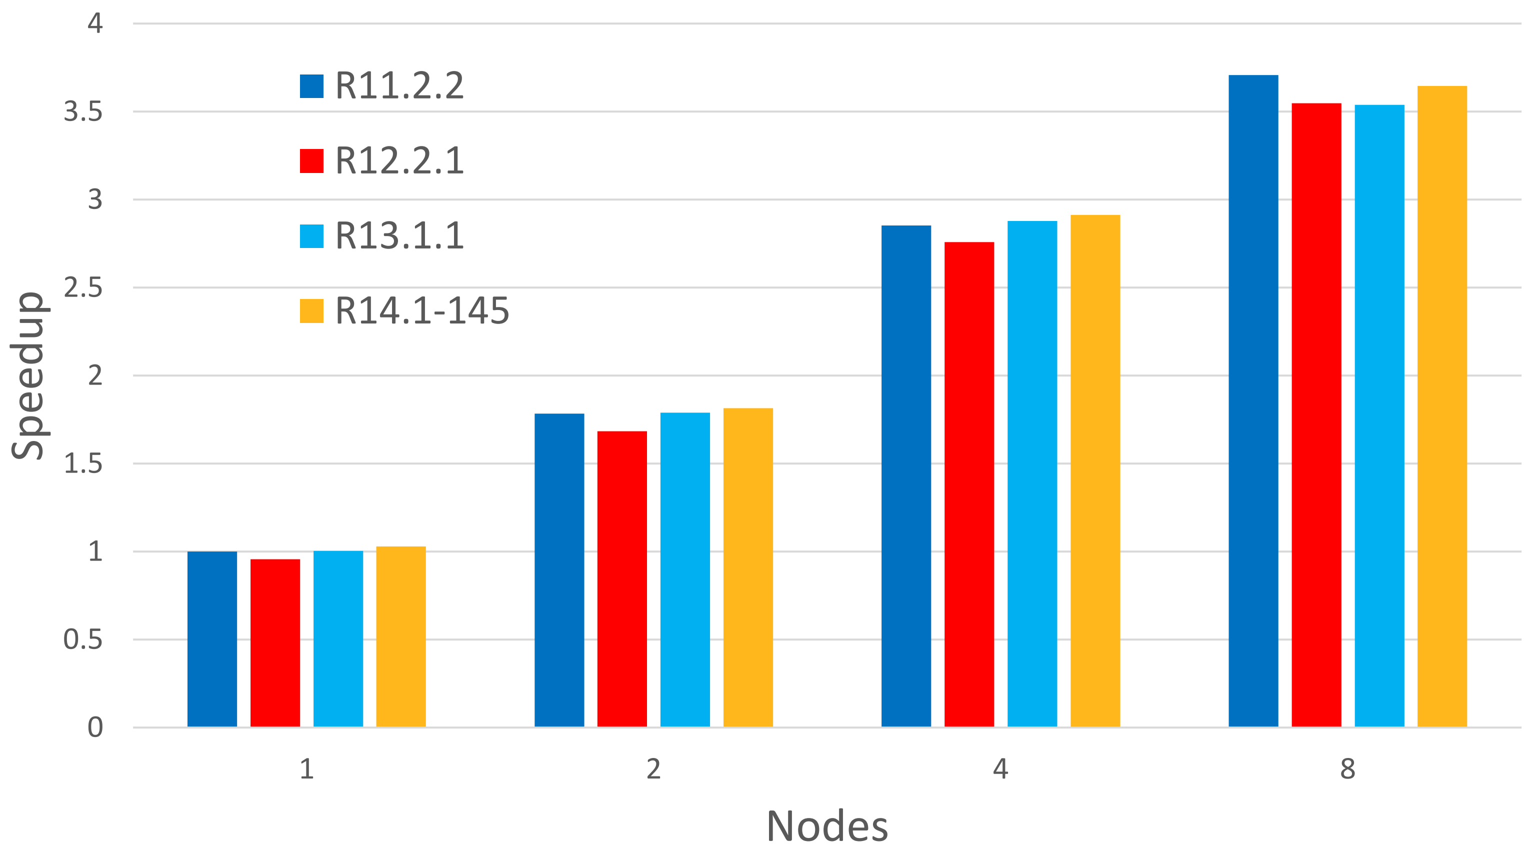

Conclusion

- Support both ifort/MKL and aocc/AOCL binaries for future releases to get best performance for target hardware
- Work with ifort and aocc development teams to avoid hardware dependent numerical noise
- X86_64 still has better computational performance than arm64. Arm64 has better price and power performance than x86-64.
- There are several next generation arm64 CPUs and we will continue support on those hardware
- There are few new CPU/GPU shared memory systems available, and we are exploring the new numerical schemes for those new hardware