

THE WORLD'S FIRST PETASCALE ARM SUPERCOMPUTER



ASTRA

"Per aspera ad astra"

VANGUARD

S.D. Hammond, K.T. Pedretti, J.H. Laros, A. Younge, C.T. Vaughan, P.T. Lin, R.J. Hoekstra
Center for Computing Research, Sandia National Laboratories, NM

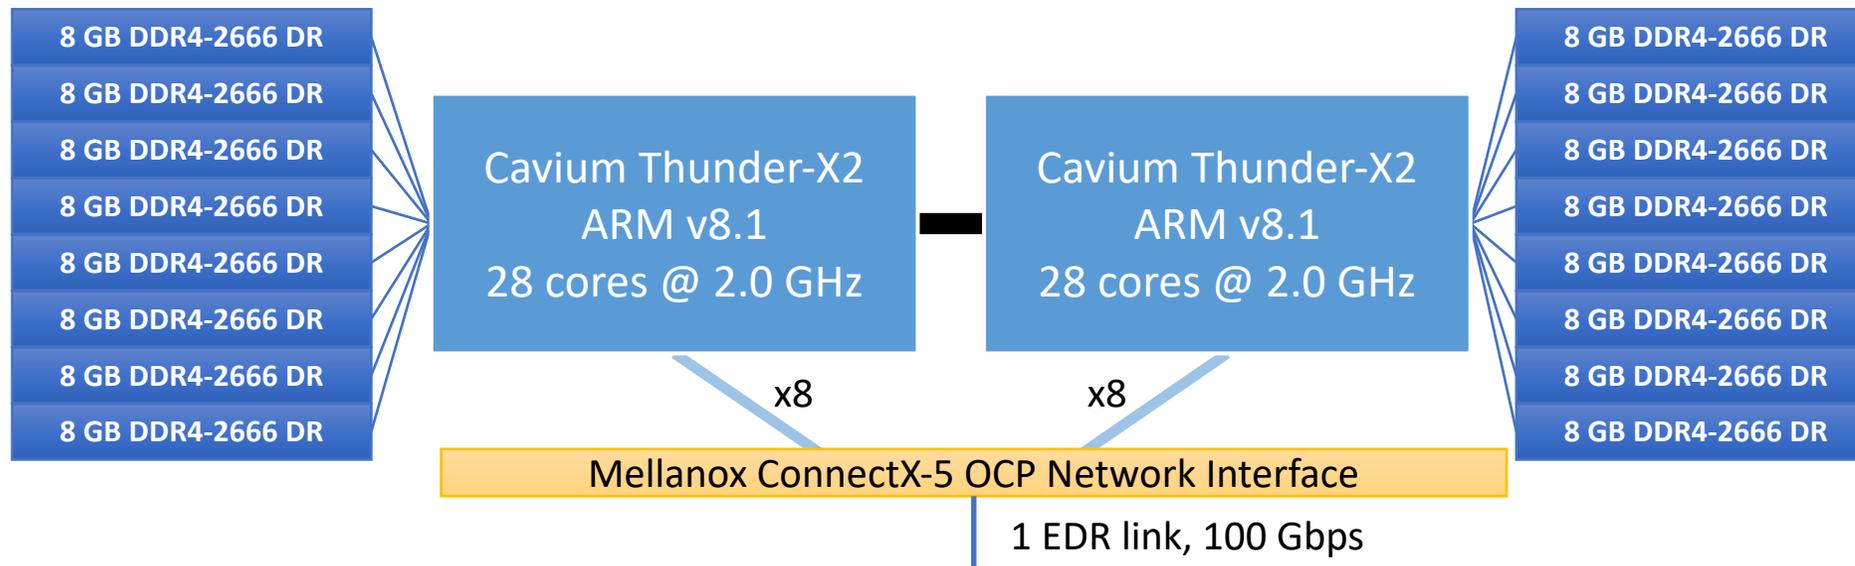
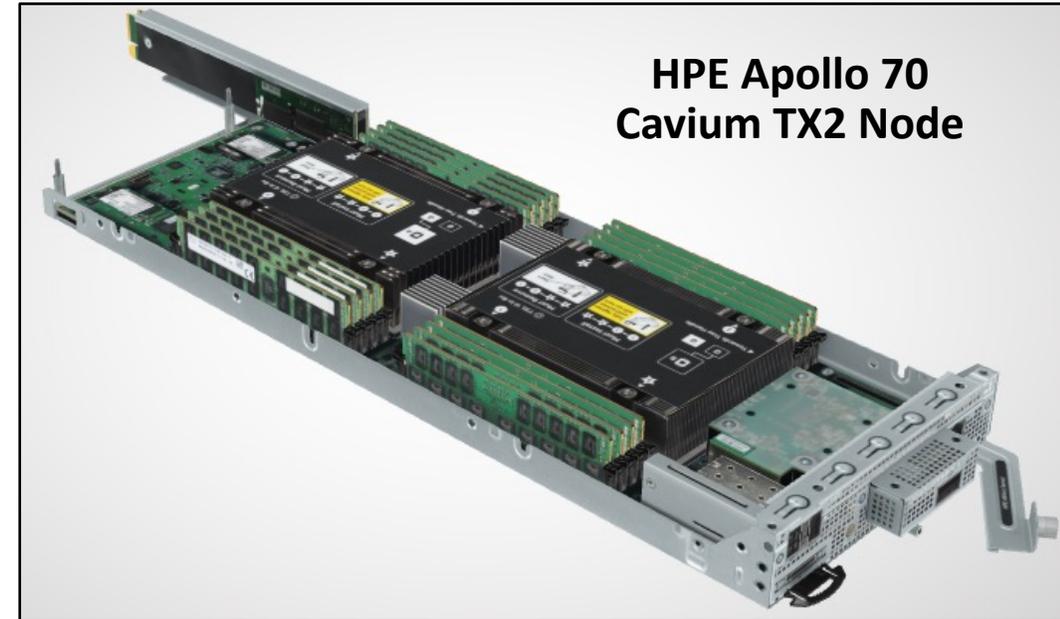
UNCLASSIFIED UNLIMITED RELEASE - SAND2019-7359 PE



Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

Astra Architecture

- **2,592** HPE Apollo 70 compute nodes
 - Cavium Thunder-X2 **Arm** SoC, 28 core, 2.0 GHz
 - 5,184 CPUs, 145,152 cores, 2.3 PFLOPs system peak
 - 128GB DDR Memory per node (**8 memory channels per socket**)
 - Aggregate capacity: 332 TB, Aggregate Bandwidth: 885 TB/s
- Mellanox IB EDR, ConnectX-5
- HPE Apollo 4520 All-flash storage, Lustre parallel file-system
 - Capacity: 403 TB (usable), recently upgraded to ~1PB
 - Bandwidth 244 GB/s

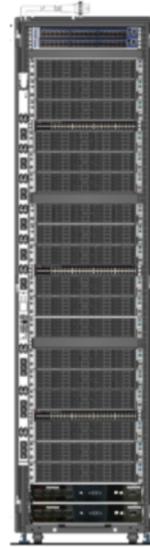


Astra – the First Petscale Arm based Supercomputer

HPE Apollo 70 Chassis: 4 nodes



HPE Apollo 70 Rack



18 chassis/rack

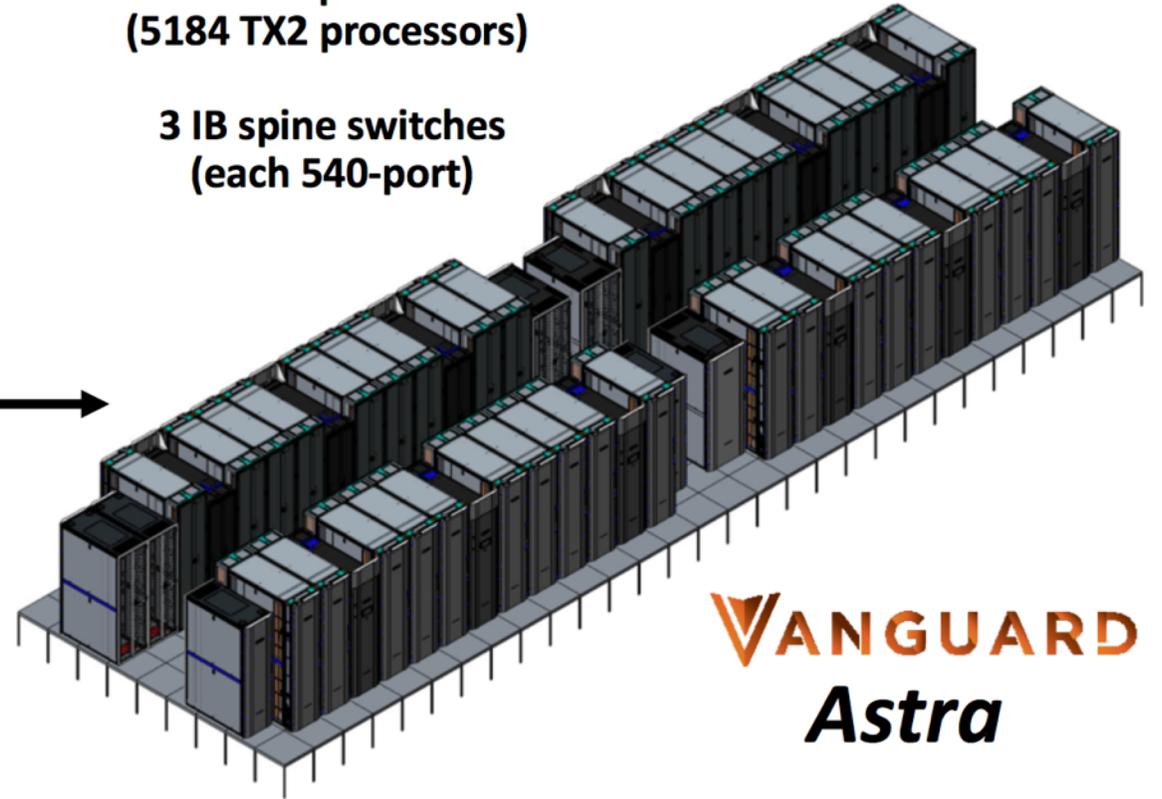
72 nodes/rack

**3 IB switches/rack
(one 36-port switch
per 6 chassis)**

**36 compute racks
(9 scalable units, each 4 racks)**

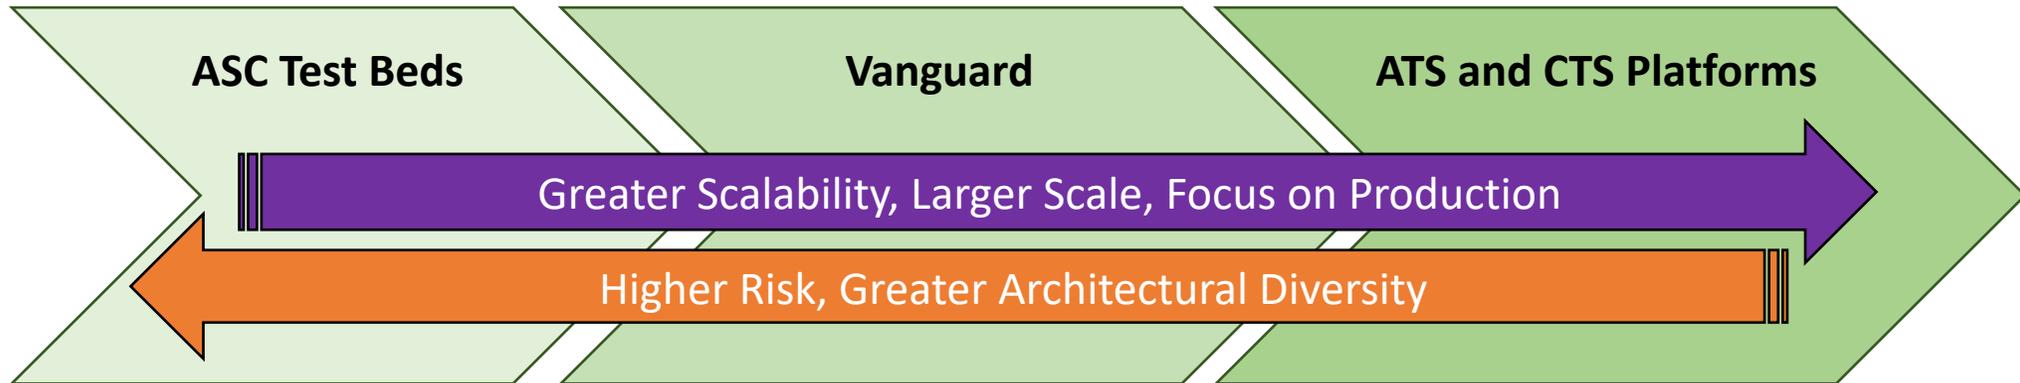
**2592 compute nodes
(5184 TX2 processors)**

**3 IB spine switches
(each 540-port)**



**VANGUARD
Astra**

Where Vanguard Fits



Test Beds

- Small testbeds (~10-100 nodes)
- Breadth of architectures Key
- Brave users

Vanguard

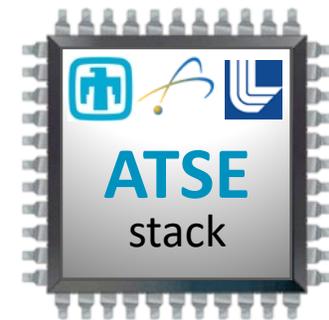
- Larger-scale experimental systems
- Focused efforts to mature new technologies
- Broader user-base
- Not Production
- **Tri-lab resource but not for ATCC runs**

ATS/CTS Platforms

- Leadership-class systems (Petascale, Exascale, ...)
- Advanced technologies, sometimes first-of-kind
- Broad user-base
- Production Use

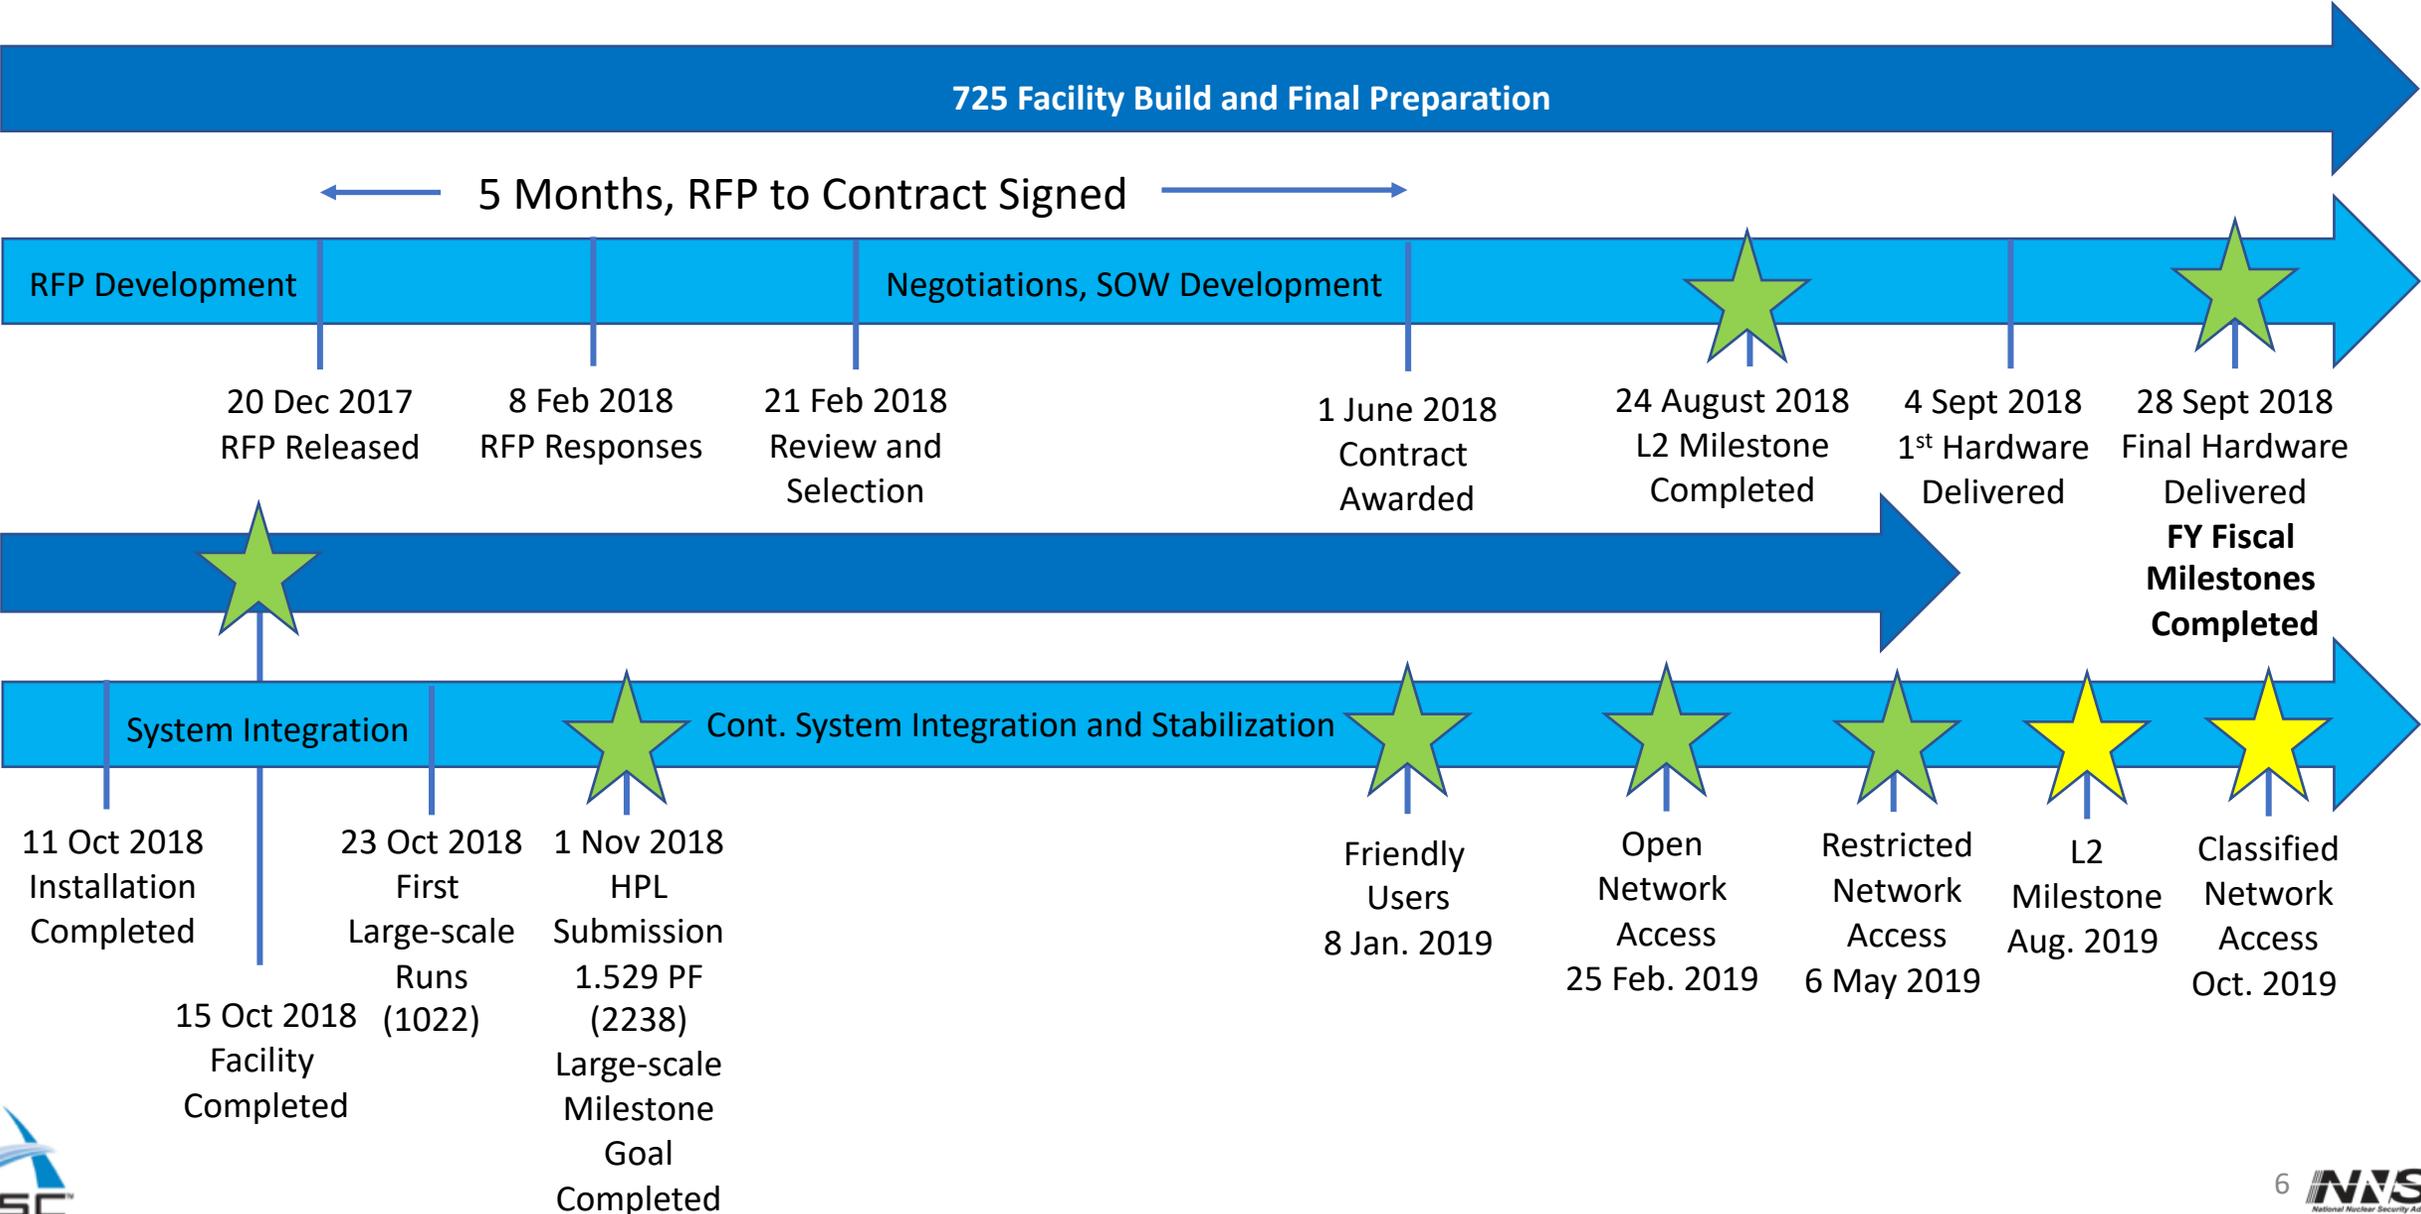
NNSA/ASC Advanced Trilab Software Environment (ATSE) Project

- Advanced Tri-lab Software Environment
 - Sandia leading development with input from Tri-lab Arm team
 - Will be the user programming environment for Vanguard-Astra
 - Partnership across the NNSA/ASC Labs and with HPE
- Lasting value
 - Documented specification of:
 - Software components needed for HPC production applications
 - How they are configured (i.e., what features and capabilities are enabled) and interact
 - User interfaces and conventions
 - Reference implementation:
 - Deployable on multiple ASC systems and architectures with common look and feel
 - Tested against real ASC workloads
 - Community inspired, focused and supported

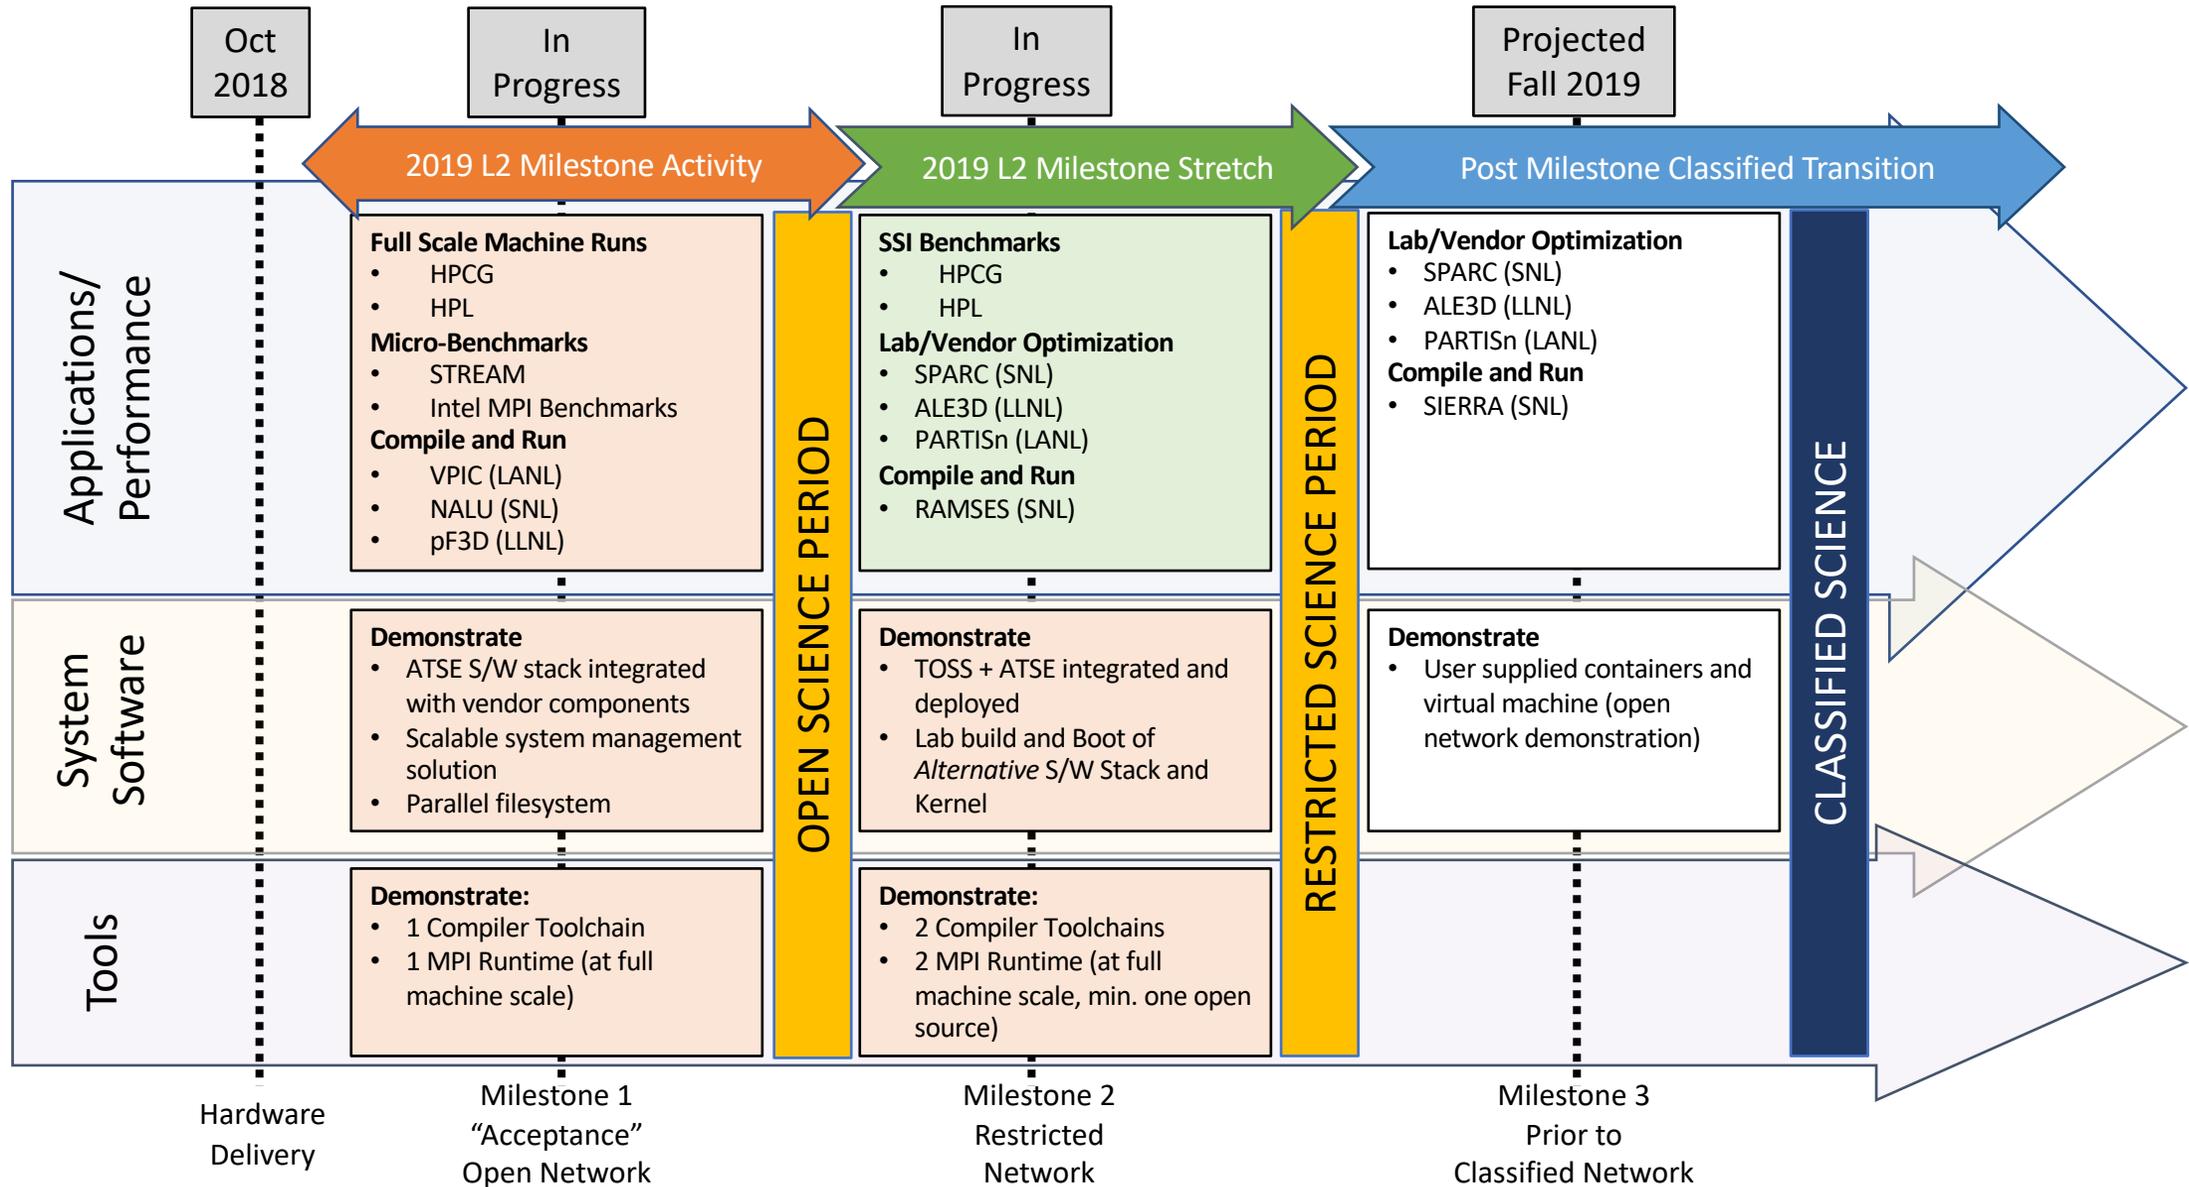


ATSE is an integrated software environment for ASC workloads

Vanguard-Astra: Timeline



Astra Milestone Schedule



Applications

Results from using Astra

STREAM Memory Bandwidth

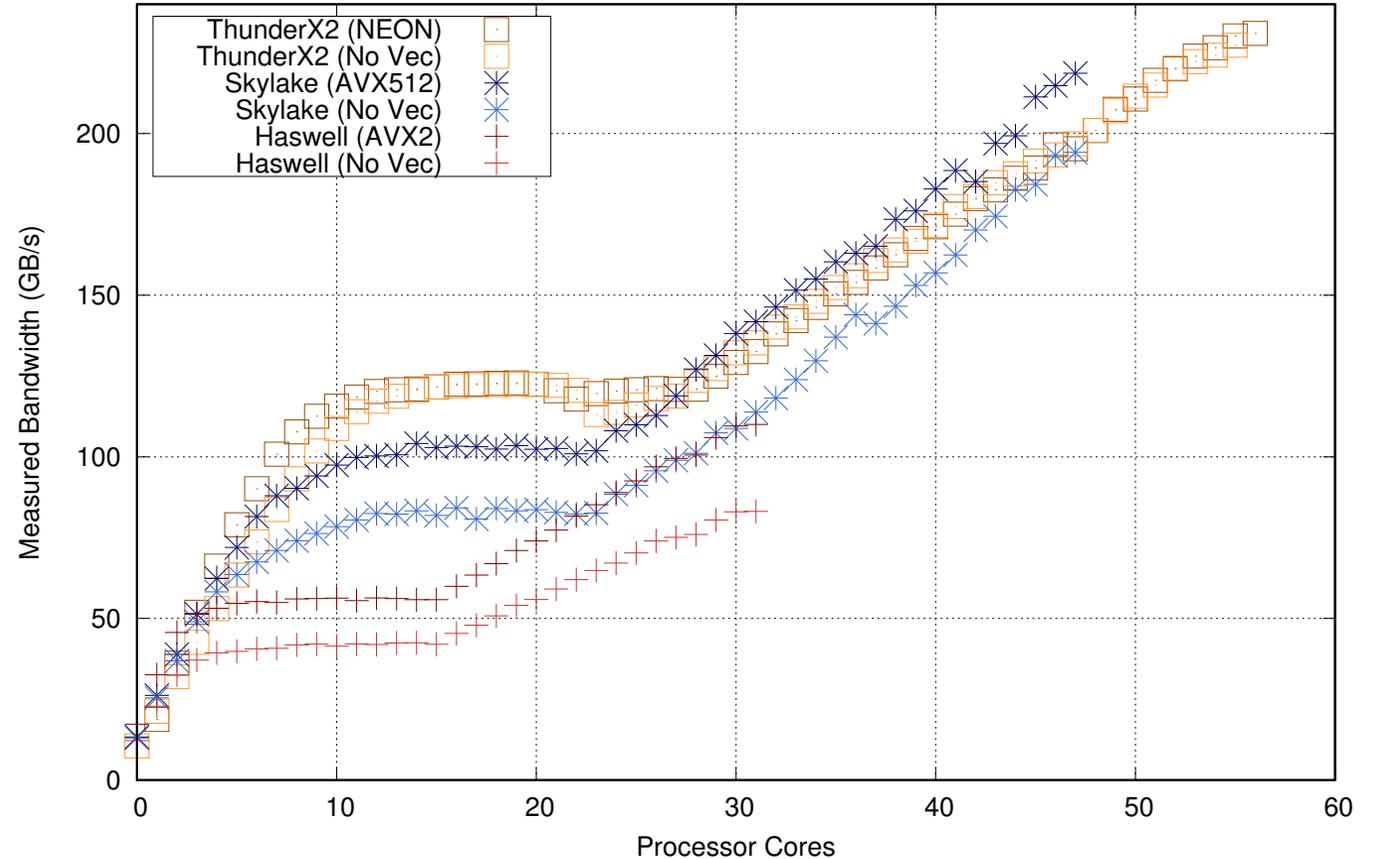
ThunderX2 provides high memory bandwidth

- 8 channels of DDR4
- Limited by ring NoC (not memory channels or DIMMs)

Vectorization makes no discernable difference to performance at large core counts

- Around 10% higher with NEON at smaller core counts (5 – 14)

Hand written kernels can exceed 250GB/s (=2X Haswell systems)



Higher is better

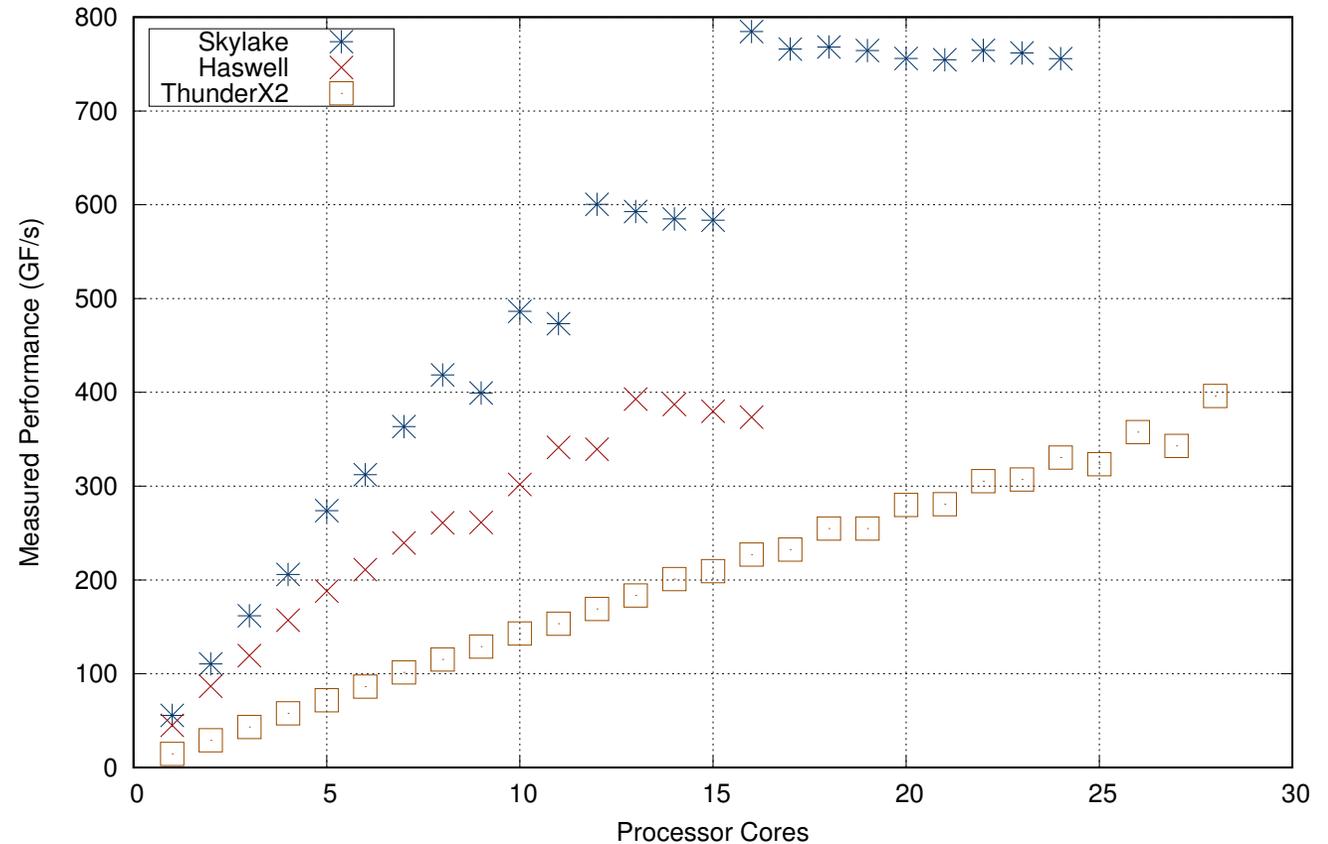
DGEMM Compute Performance

ThunderX2 has similar performance at scale to Haswell

- Roughly twice as many cores (TX2 has 1.75X cores)
- Half the vector width (TX2 = 2 vs. HSW = 4, SKX = 8)

See strata in Intel MKL results, usually a result of matrix-size kernel optimization

- ARM PL provides smoother performance results (essentially linear growth)



Higher is better

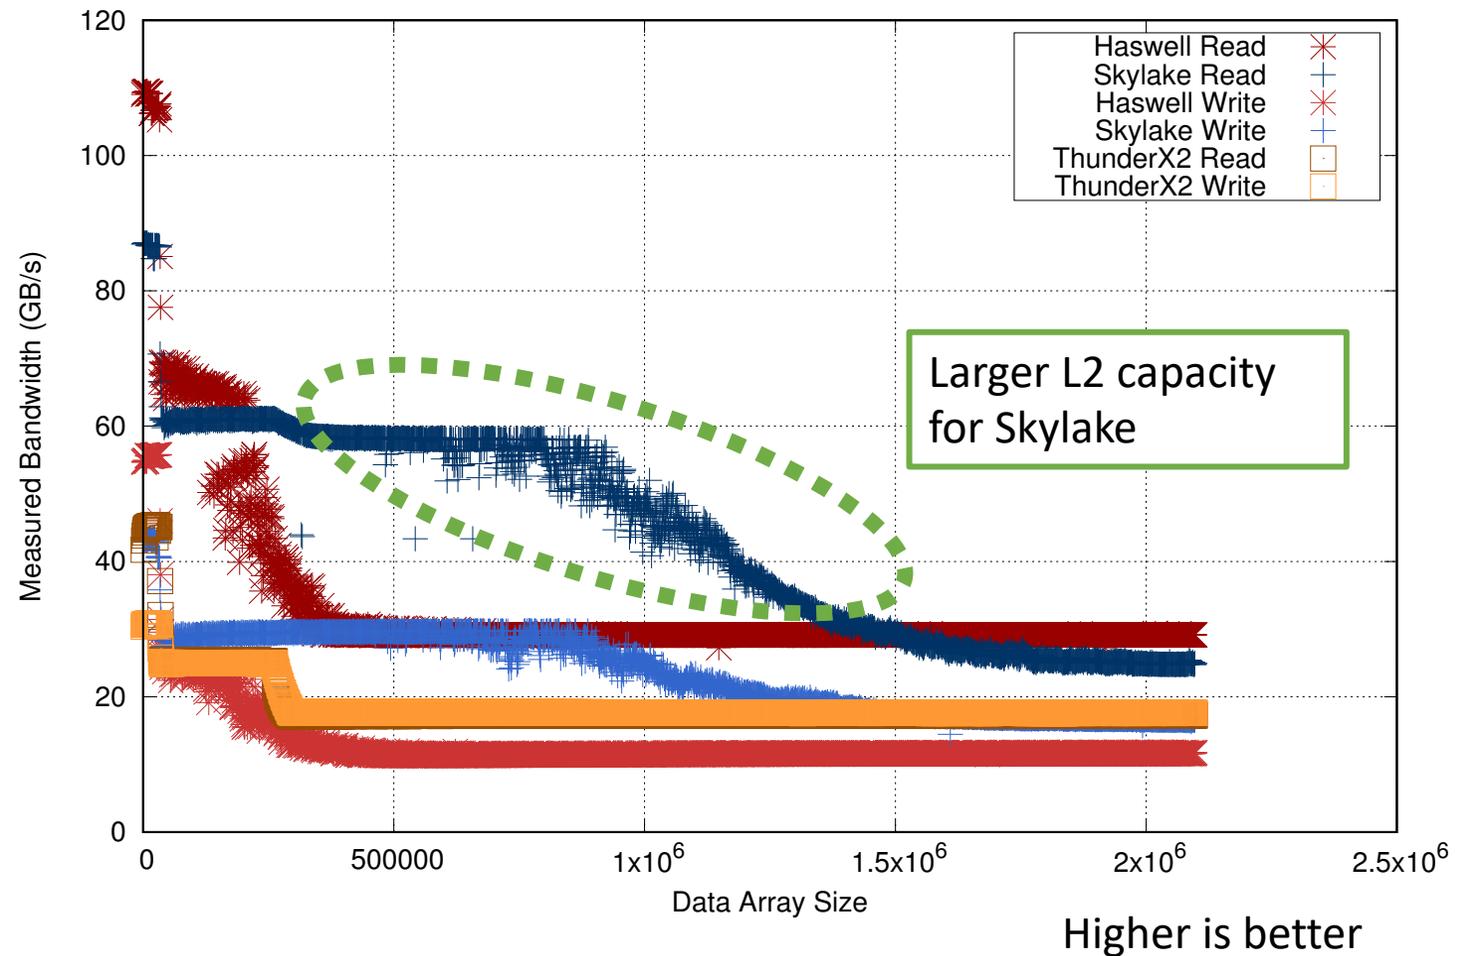
Cache Performance

Haswell has highest per-core bandwidth (read and write) at L1, slower at L2

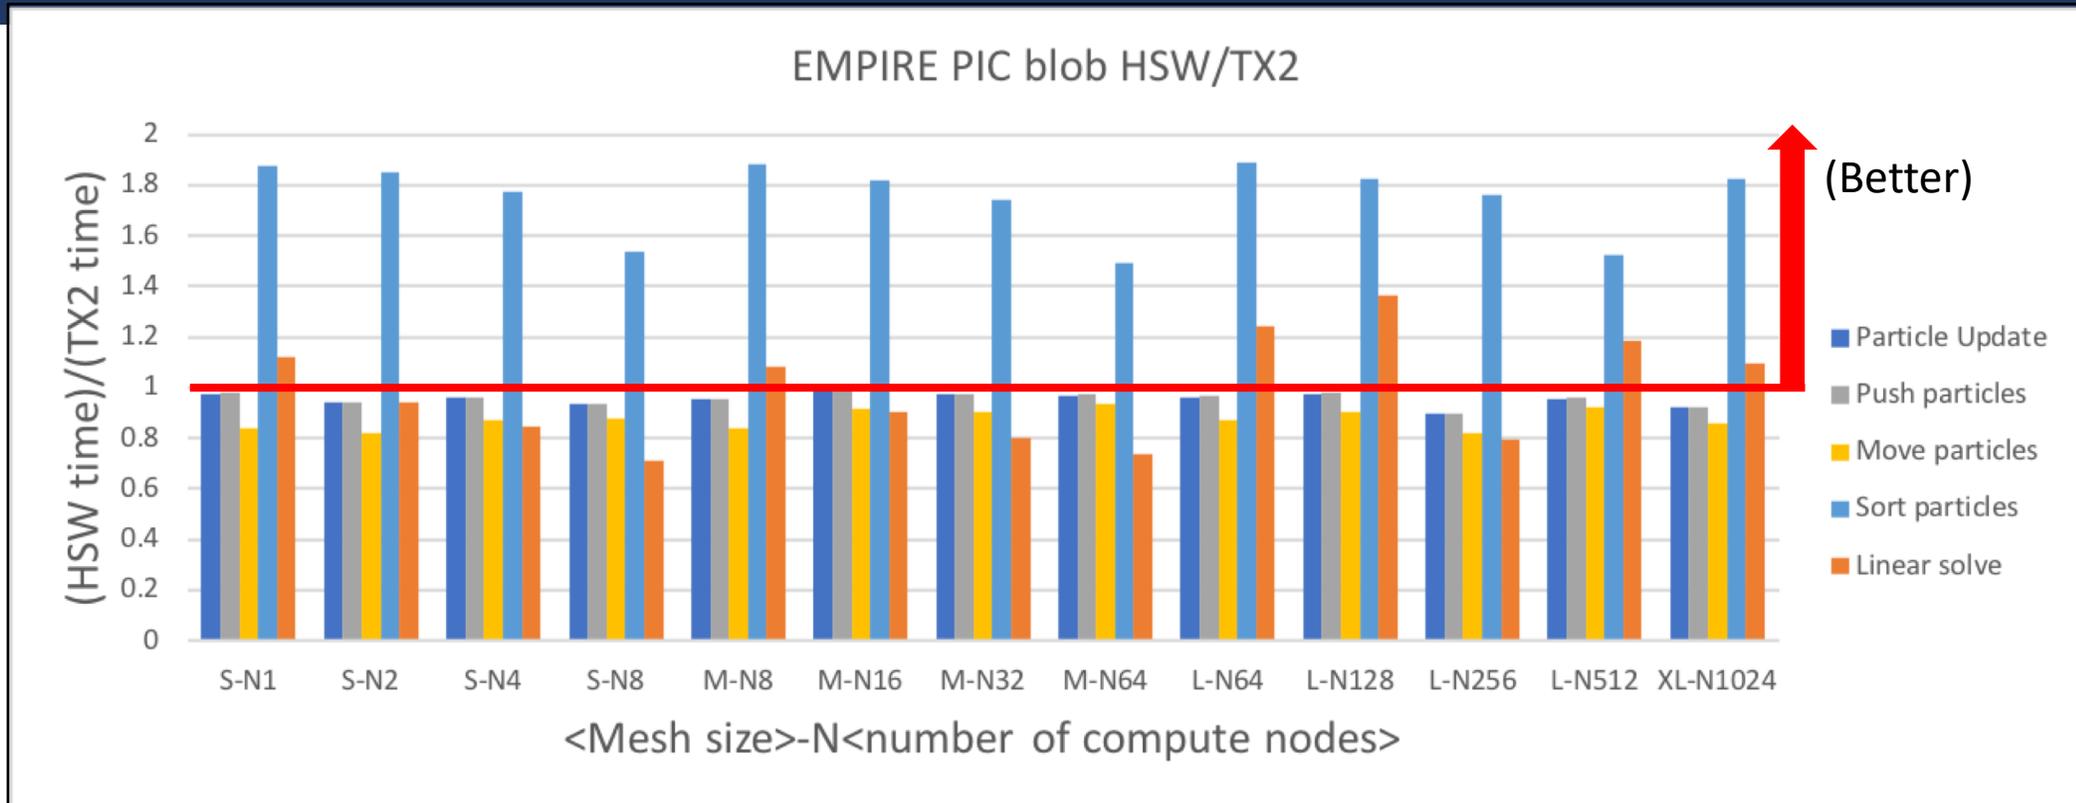
Skylake redesigned cache sizes (larger L2, smaller L3) shows up in graph

- Higher performance for certain work-set sizes (typical for unstructured codes)

TX2 more uniform bandwidth at larger scale (see less asymmetry between read/write)



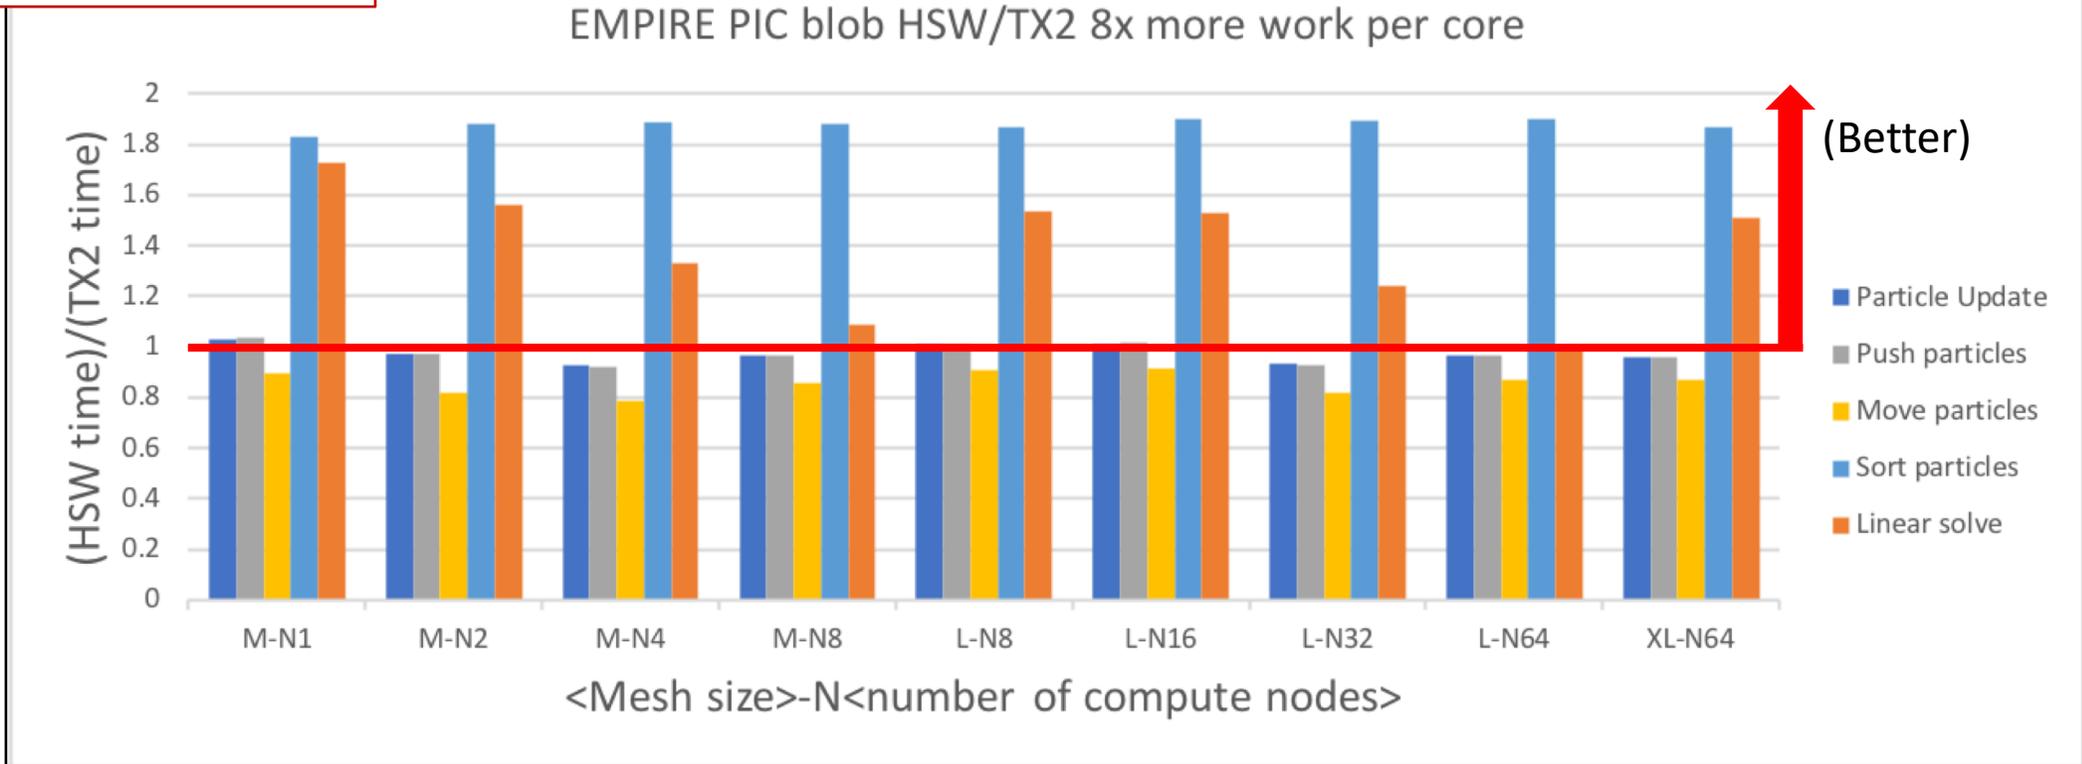
EMPIRE on Astra



- TX2 node has ~2x memory bandwidth and 1.75x cores (56 vs. 32) of Trinity HSW node
- $(\text{HSW time})/(\text{TX2 time}) > 1$ means TX2 is faster
- Strong scaling for awesome blob small mesh (1-8 nodes), strong scaling for medium mesh (8-64 nodes), strong scaling for large mesh (64-512)
- $(\text{HSW time})/(\text{TX2 time})$ for linear solve not great, low computation/communication regime

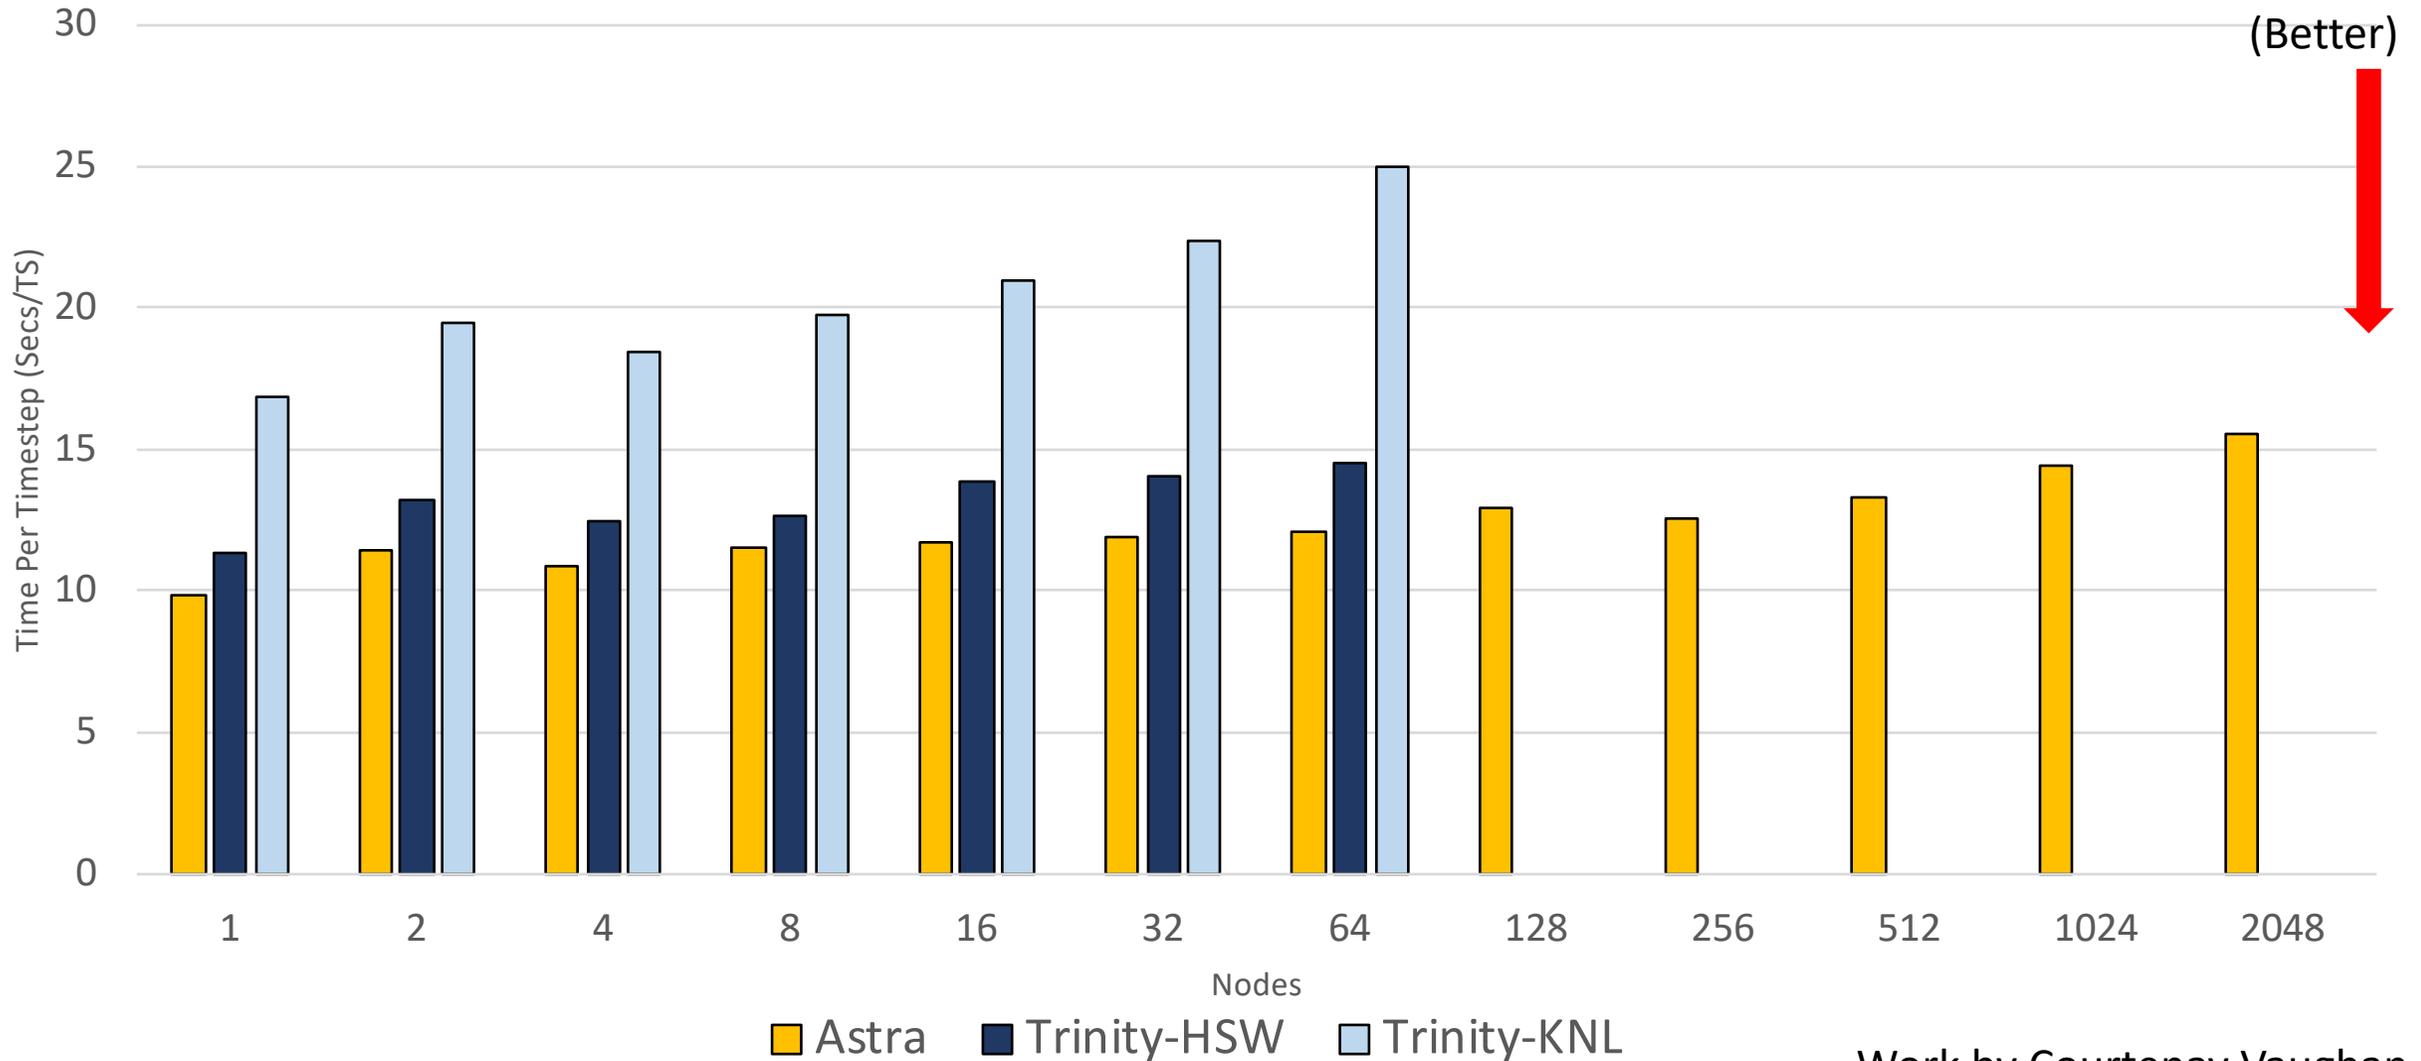
Larger Problems gain more from memory B/W

EMPIRE on Astra

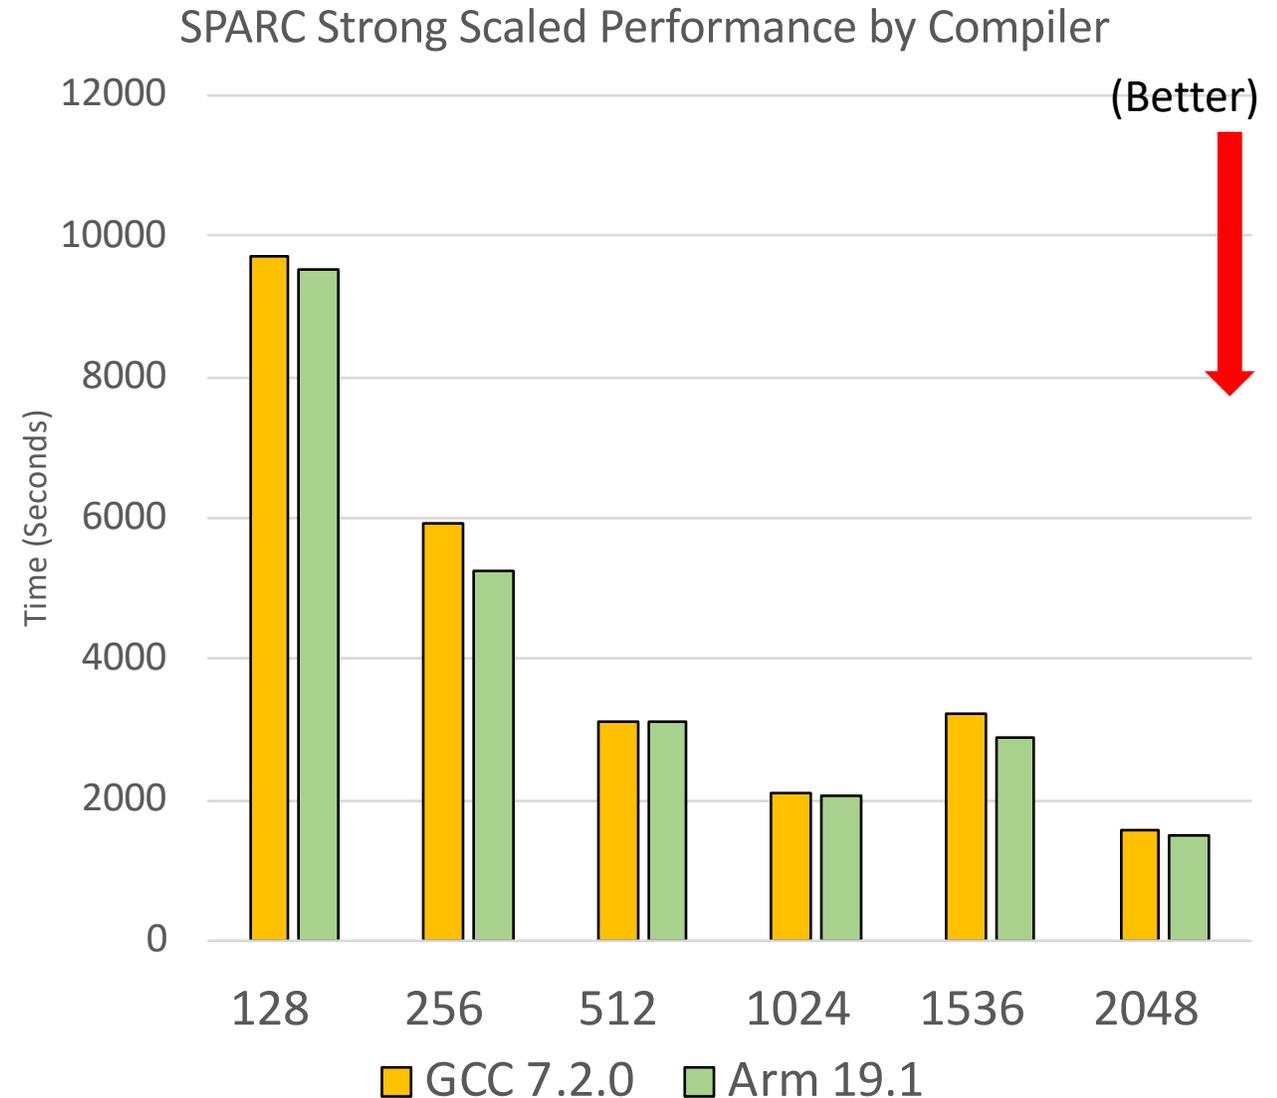


- TX2 node has ~2x memory bandwidth and 1.75x cores (56 vs. 32) of Trinity HSW node
- $(\text{HSW time})/(\text{TX2 time}) > 1$ means TX2 is faster
- Strong scaling for awesome blob small mesh (1-8 nodes), strong scaling for medium mesh (8-64 nodes), strong scaling for large mesh (64-512)
- $(\text{HSW time})/(\text{TX2 time})$ for linear solve not great, low computation/communication regime

Shape Charge Problem (Mixed-Material, Benchmark Problem, Fortran)



- SPARC is Sandia's latest CFD modeling code
 - Developed under NNSA ATDM Program
 - Written to be threaded and vectorized
 - Uses Kokkos programming abstractions
 - Approximately 2-3M lines of code for binary (including Trilinos packages, mostly C++, tiny bit of Fortran)
- Mixture of assembly and solve phases
- Successfully compiles with GCC and Arm HPC compilers on Astra
- Results show performance with Arm HPC compiler varies from 0.5% faster than GCC to 10% slower
 - This seems to be consistent across our code portfolio at present



- NALU – Large Scale CFD Simulation
 - Used as a proxy for large-scale engineering code suite
 - Same mesh handling (I/O and distribution) library
 - Drives Trilinos solvers using multi-grid libraries
- Results show strong solve kernel performance but slower assembly
 - Some routines do not scale well with increasing MPI rank counts (problem on Astra and KNL)

NALU Timesteps per 24 Hours @ 2048 Nodes

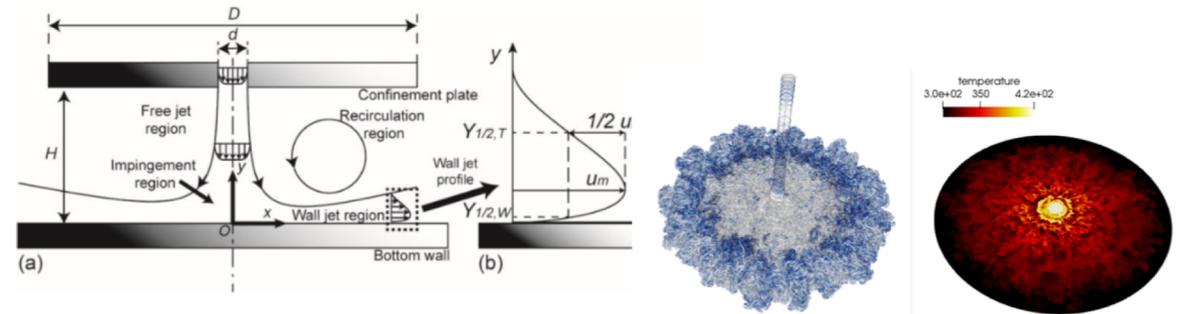
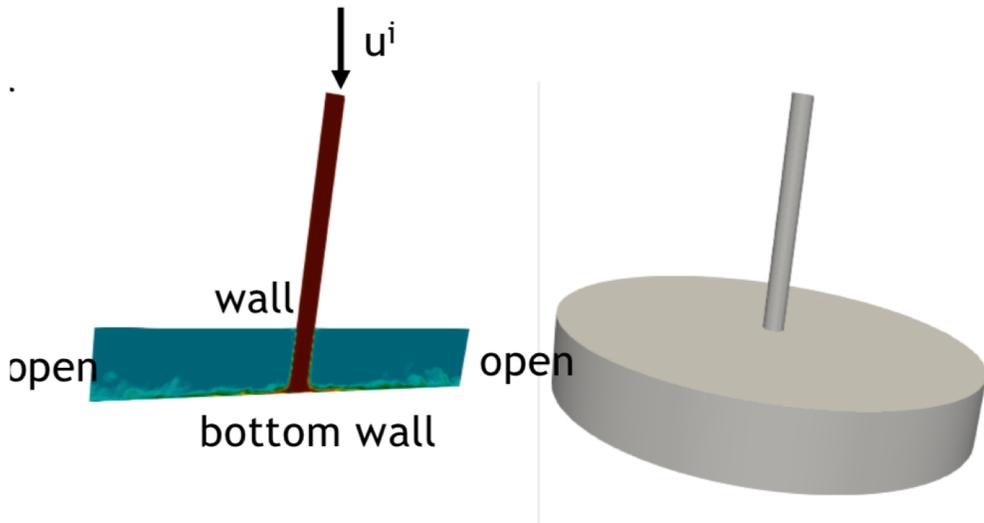
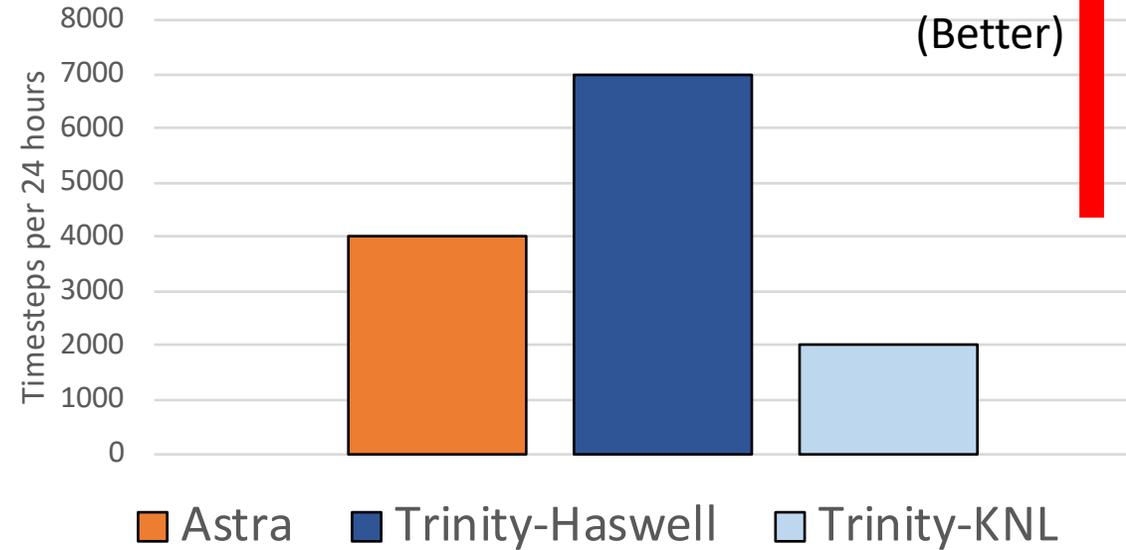
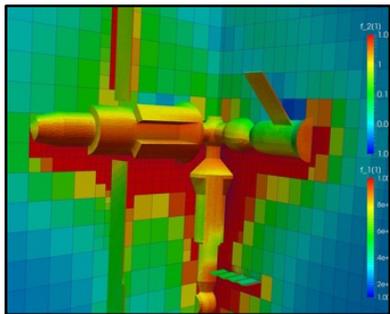


Figure 1. Schematic drawings of (a) the axisymmetric flow field formed by the impinging jet and, (b) the wall jet structure and nomenclature.

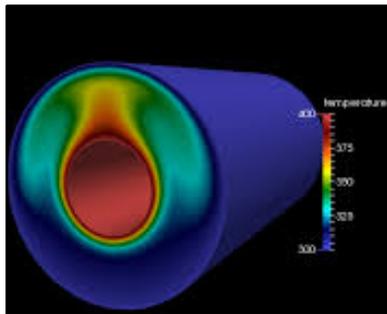
Early Results from Astra

- ThunderX2 is less reliant on vectorization to utilize available memory bandwidth.
 - Cores can consume available memory bandwidth without vectorized code.
 - Downside: vector units are small so compute-dense code may run slower, extra cores help offset this when comparing node-to-node
- Most of our complex solver libraries and applications compile with GCC or Arm compilers without significant issues.
 - Functional portability for broad code portfolio without significant code rework (NALU, SPARC, CTH, etc.)
 - Acid test is getting the performance out of generated code
- Cache performance will likely impact some of our codes that have reasonable locality
 - Suspect that caches simply perform slower on TX2 versus Xeon
 - Lack of support for gather operations
- Most packages ported and running on the platform, ATSE environment has worked out well



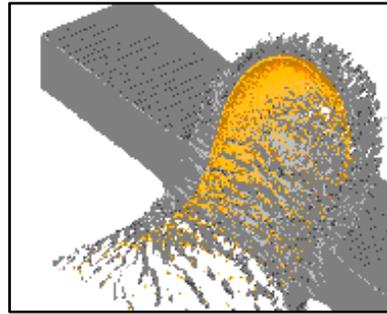
Monte Carlo

1.60X



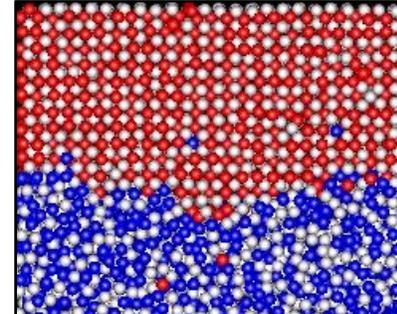
CFD Models

1.45X



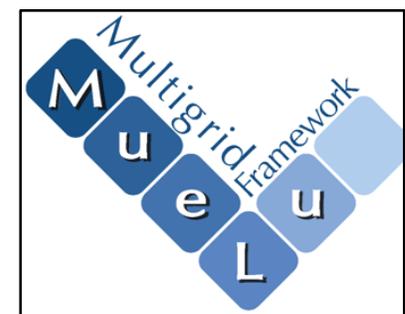
Hydrodynamics

1.30X



Molecular Dynamics

1.42X

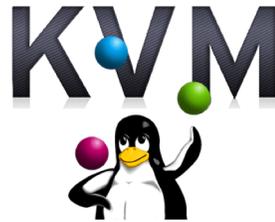
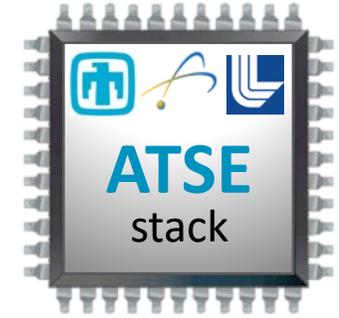


Linear Solvers

1.87X

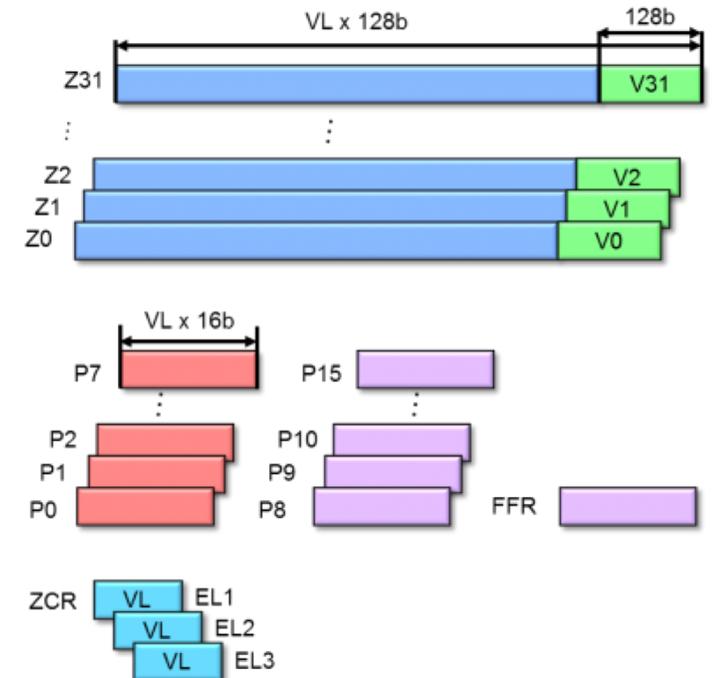
ATSE R&D Efforts – Developing Next-Generation NNSA Workflows

- Workflows leveraging containers and virtual machines
 - Support for machine learning frameworks
 - ARMv8.1 includes new virtualization extensions, SR-IOV
- Evaluating parallel filesystems + I/O systems @ scale
 - GlusterFS, Ceph, BeeGFS, Sandia Data Warehouse, ...
- Resilience studies over Astra lifetime
- Improved MPI thread support, matching acceleration
- OS optimizations for HPC @ scale
 - Exploring spectrum from stock distro Linux kernel to HPC-tuned Linux kernels to non-Linux lightweight kernels and multi-kernels
 - Arm-specific optimizations



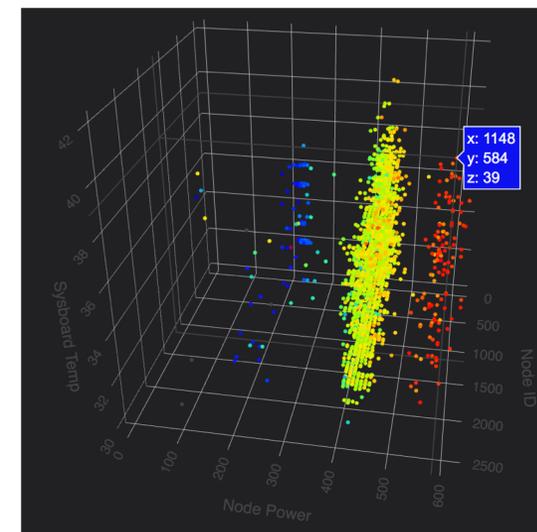
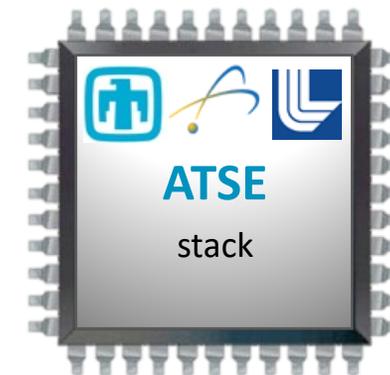
SVE Enablement

- SVE work is underway
 - Using ArmIE (fast emulation) and RIKEN GEM5 Simulator
 - GCC and Arm toolchains
- Collaboration with RIKEN
 - Visited Sandia (participants from SNL, LANL, LLNL, RIKEN)
 - Discussion of performance and simulation techniques
 - Deep-dive on SVE (GEM5)
- Short term plan
 - Use of SVE intrinsics for Kokkos-Kernels SIMD C++/data parallel types
 - Underpins number of key performance routines for Trilinos libraries
 - Seen large (6X) speedups for AVX512 on KNL and Skylake
 - Expect to see similar gains for SVE vector units
 - Critical performance enablement for Sandia production codes



Vanguard Astra: Lessons Learned *or* Reasons to Prototype new Technologies

- Similar to CTS, NALU and other applications were forcing out of spec voltage swings
 - In this case memory bus
- Pesky fabric instability
 - Lots of hands on nodes probable cause
- Pioneering new systems management solution (HPCM) with vendor
 - Combined with new software stack (ATSE)
- At scale testing reveals previously unseen issues
 - kworker bug not seen on Comanche or x86 platforms
- Systems monitoring is CRITICAL (debug/analyze many of above)
- Early hardware requires frequent and quick iterations of software stack
 - Tension with accelerated move to classified where this is a challenge
- Keeping system in sync (updating software) a challenge -> future work with containers



Summary

- Astra deployed and applications work is now well underway
 - Sandia has main ATDM application portfolio and libraries ported
 - Application scaling up to 2,048 nodes routinely running
- Performance Assessment:
 - **Depends on your code** – similar to Haswell for cache-bound/compute-bound, faster than Skylake for memory bandwidth. In reality most costs are a mixture of these
 - Arm HPC compilers are functional, need work to improve code performance (NNSA/ASC Arm Compiler contract)
- Planning for migration to classified network is well underway, expected in Fall 2019

It Takes an Incredible Team...

- DOE Headquarters:

- Thuc Hoang
- Mark Anderson

- Sandia Procurement

- Sandia Facilities

- Colleagues at LLNL and LANL

- Trent D'Hooge
- Mike Lang
- Rob Neely
- Dave Rich
- Matt Leininger

- Incredible team at Sandia

- HPE:

- Mike V. and Nic Dube
- Andy Warner
- John D'Arcy
- Steve Cruso
- Lori Gilbertson
- Cheng Liao
- John Baron
- Kevin Jamieson
- Tim Wilcox
- Charles Hanna
- Michael Craig
- And loads more ...

- Cavium/Marvel:

- Giri Chukkapalli (now NVIDIA)
- Todd Cunningham
- Larry Wikelius
- Kiet Tran
- Joel James
- And loads more...

- ARM:

- ARM Research Team!
- ARM Compiler Team!
- ARM Math Libraries!
- And loads more...



Exceptional Service in the National Interest

<http://vanguard.sandia.gov>