SAND2021-1622PE

SPARC Portability and Performance

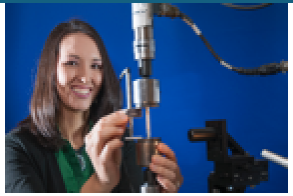

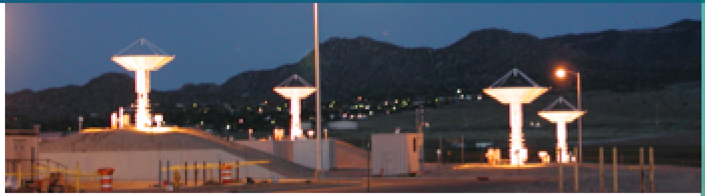

Victor Brunini

Presenting for the SPARC team

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & 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.

SPARC Portability & Performance Outline

1. Portability Strategy
2. Performance Introduction
3. SPARC Performance Results
4. Summary

SPARC Performance Portability Strategy

Leverage Trilinos for performance-portable linear solvers and I/O facilities

Implement domain-specific data structures & mesh iteration abstractions on top of Kokkos

- Can tune implementation for different platforms to maximize performance
  - Atomics vs graph coloring
  - Memory layouts & iteration patterns
  - SIMD types

All physics code remains platform agnostic

96.1% of SPARC code base is platform agnostic

L1 Performance Analysis Use-Case

Steady-state aero, 5sp/2T, RANS (11 dofs/cell), representative of analysis usage

Evaluate performance on:

1. CTS-1 (Eclipse) Xeon-Broadwell (BDW)
2. ATS-1 (Trinity) Knight's Landing (KNL)
3. ATS-2 (Sierra) Power9/V100 (V100)
4. Vanguard-1 (Astra) ARM Thunder X2 (TX2)

Four levels of mesh refinement:

Mesh	# of cells
R0	4.2M
R1	33.5M
R2	268.4M
R3	2.2B

Relative Node Performance (measured against CTS-1 systems)

		CTS1	ATS-1/Trinity	ATS-2/Sierra	Astra
		Broadwell	KNL	V100 GPU	ThunderX2
LINPACK FLOP Rates (per Node)	Perf	1.09 TF/s	~2.06 TF/s	~21.91 TF/s	~0.71 TF/s
	Rel	1.00X	1.89X	20.01X	0.65X
Memory Bandwidth (STREAM) (per Node)	Perf	~136 GB/s	~90 GB/s - ~350 GB/s	~3.4 TB/s	~250 GB/s
	Rel	1.00X	0.66X – 2.50X	25.00X	1.84X
Power (TDP, per Node)	Perf	240W	~250W	1.2kW	360W
	Rel	1.00X	1.04X	5.00X	1.50X

Generic RV Performance – Strong Scaling, Overall Runtime

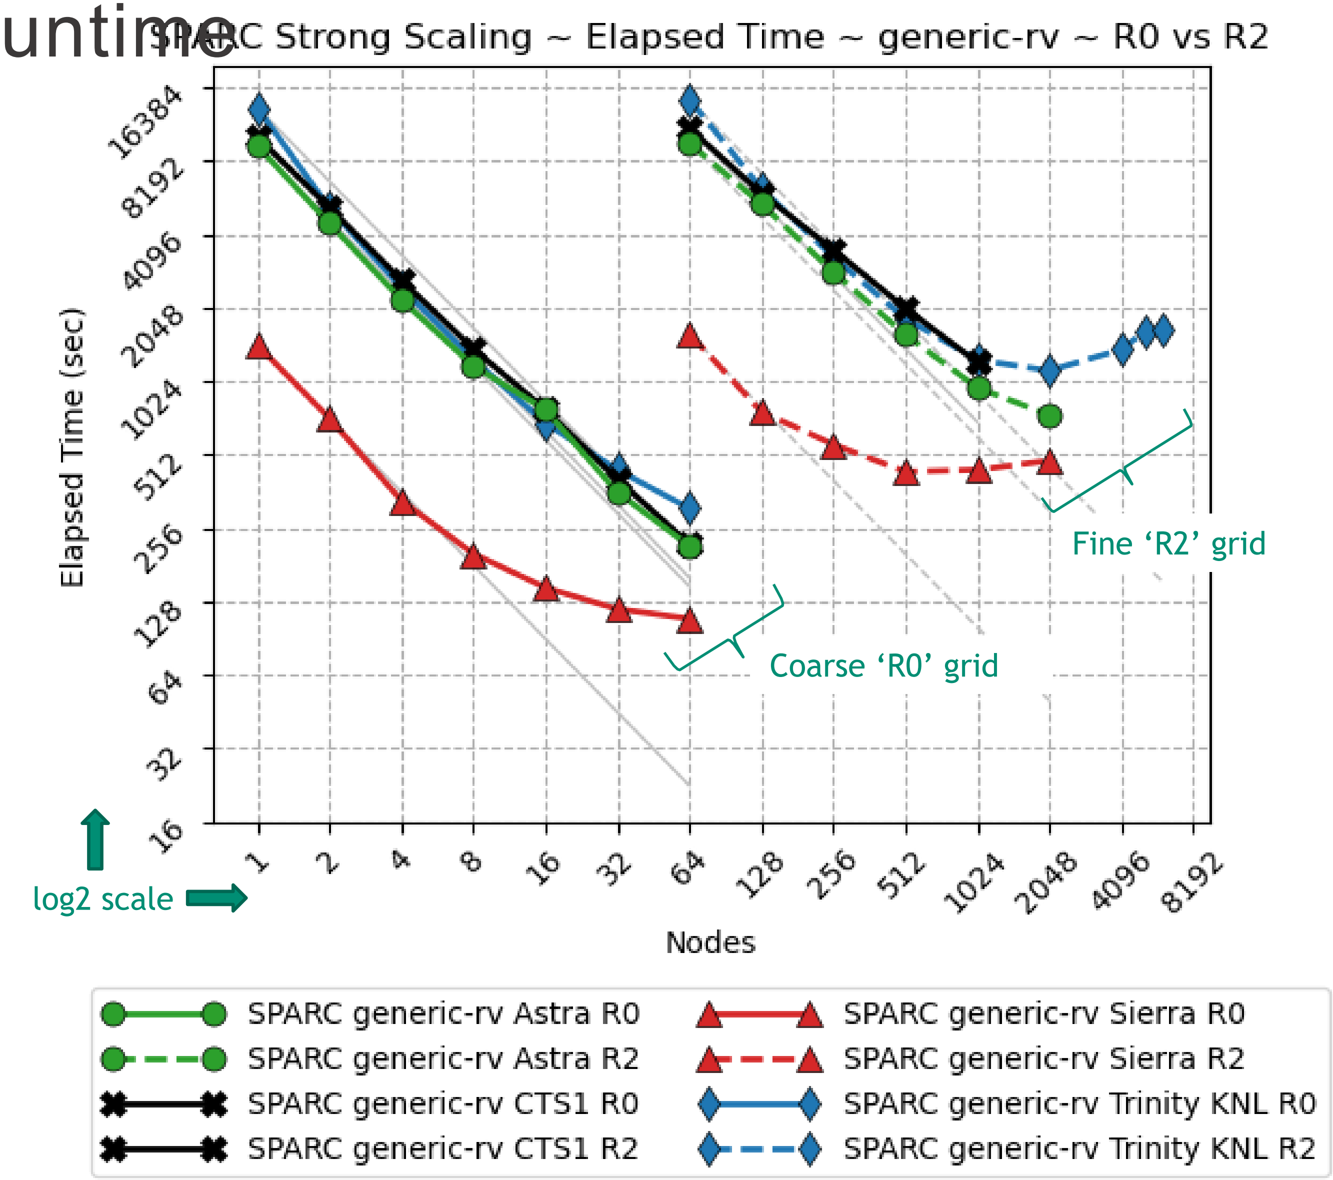

CTS1 as baseline:
Trinity: 0.8-1.2x
Astra: 1.2-1.4x
Sierra: 2-8x

Some increase in elapsed time at largest scales for Trinity and Sierra.

Generic RV Performance – Strong Scaling, Time per Time Step

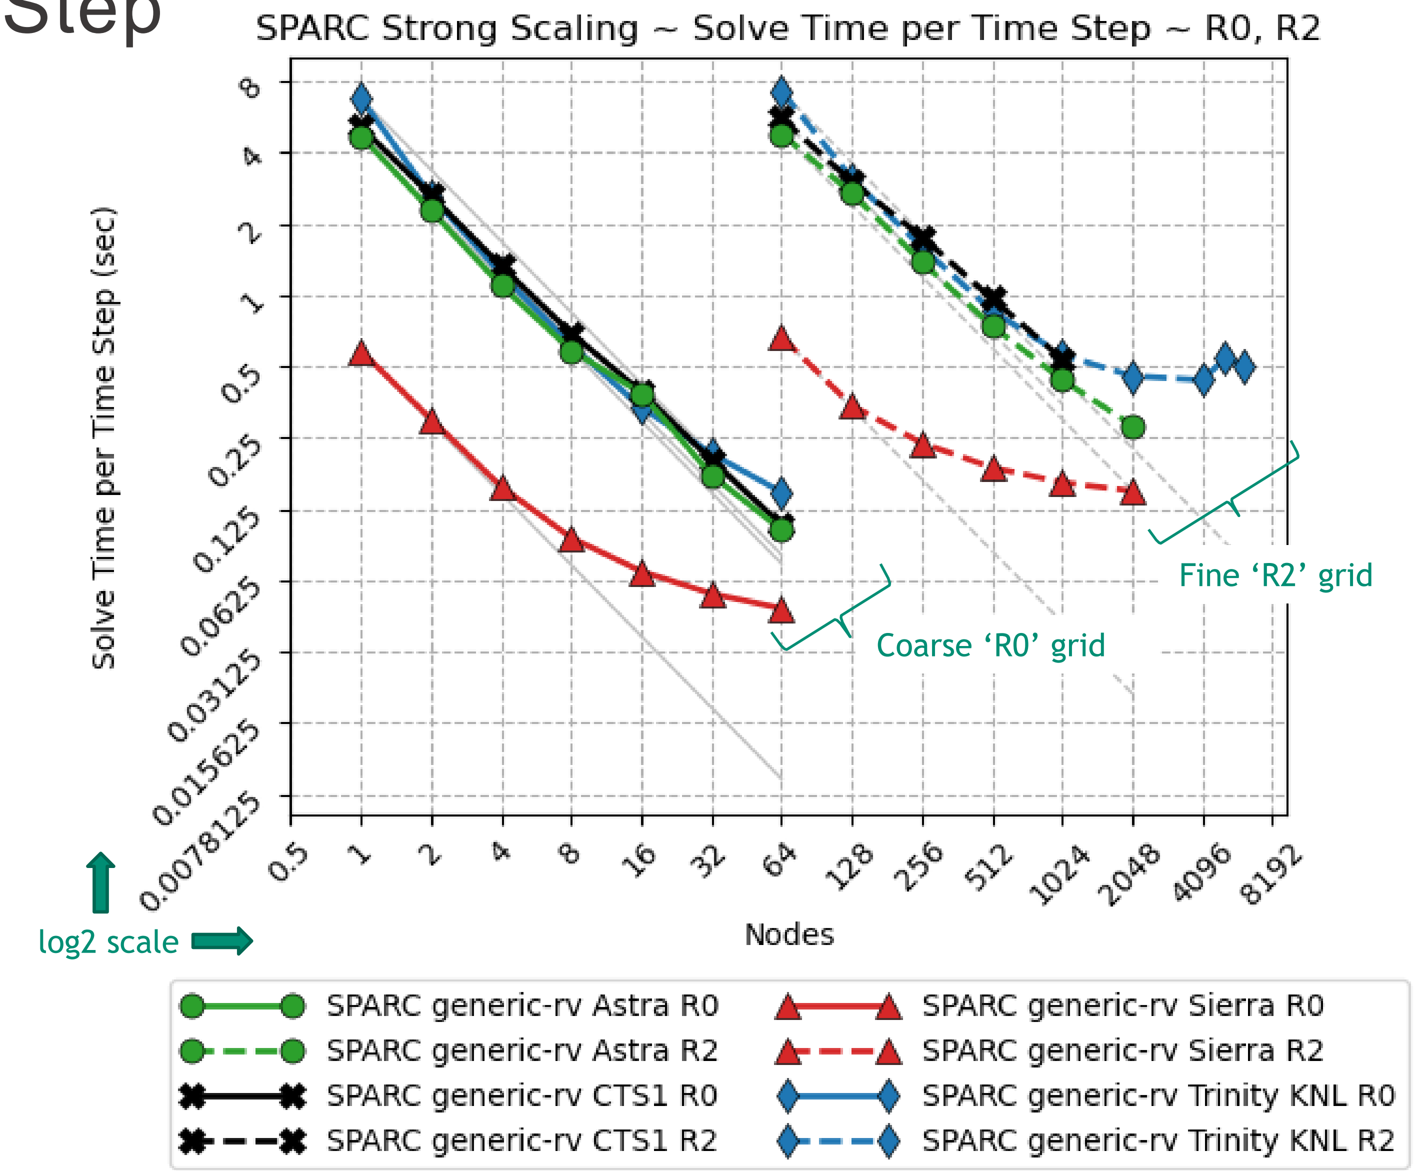

Time per time step does not show same increase at large scale on Trinity and Sierra.

Generic RV Performance – Strong Scaling - Initialization

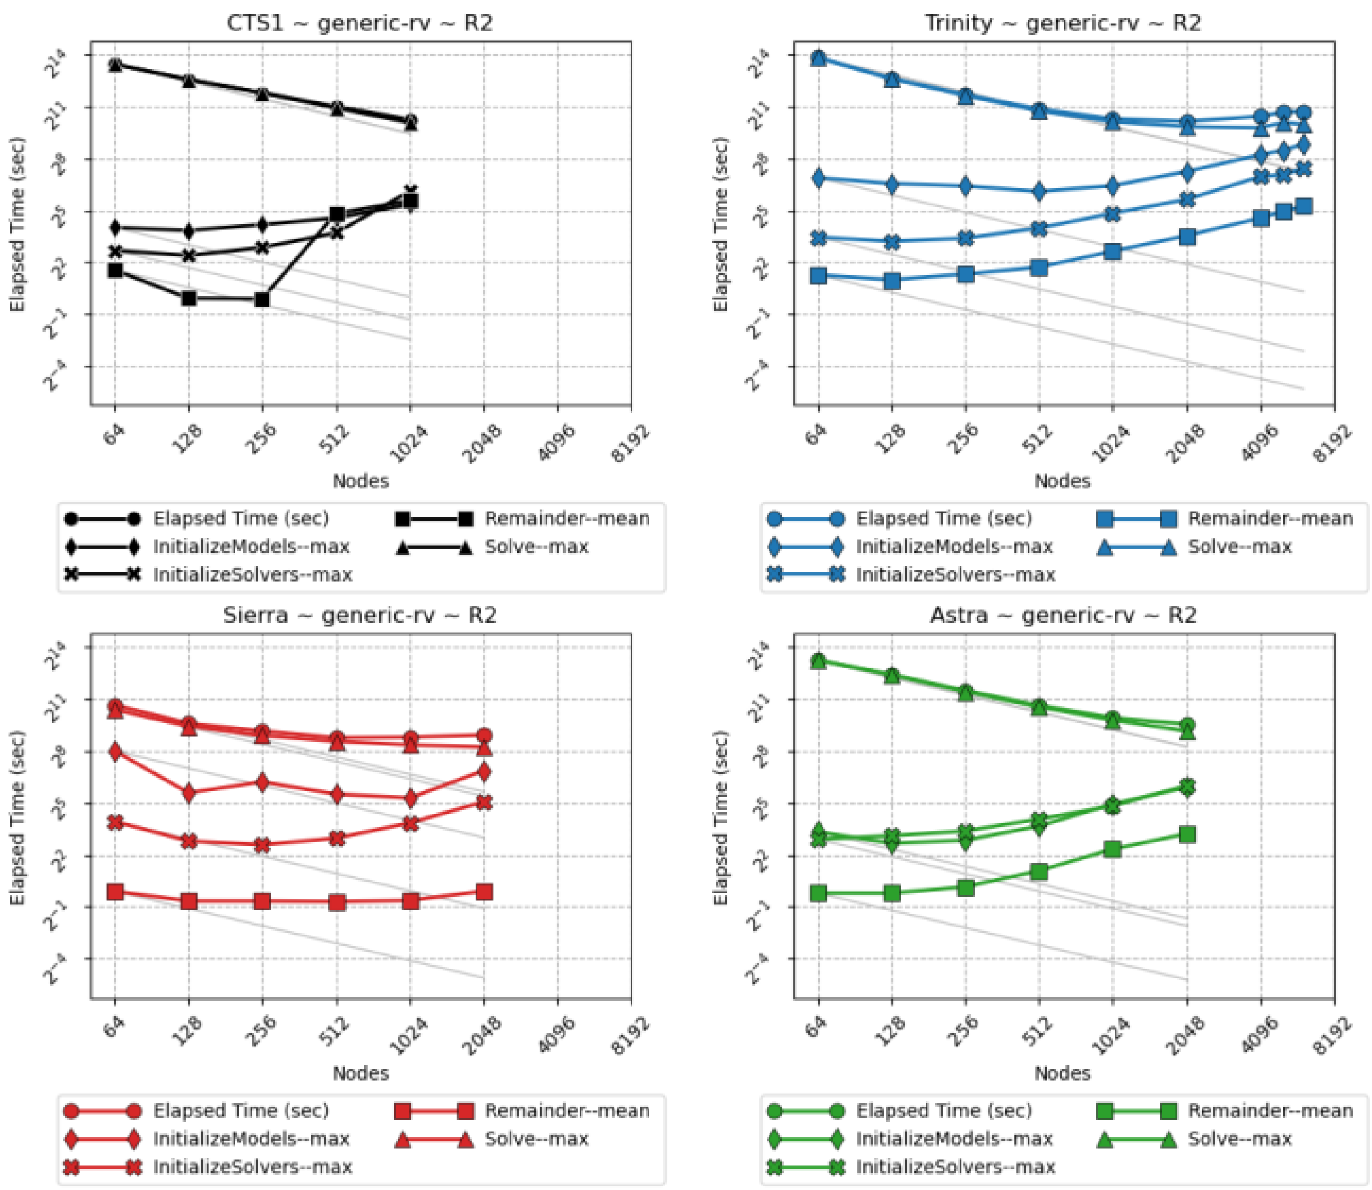

Scalability issues in initial IO setup drive up-turn in Trinity & Sierra strong scaling.

Understanding Sierra Strong Scaling

SPARC's block tridiagonal linear solver is the primary limiter of strong scaling.

1. Relative cost of compute kernels is much lower than on other platforms
2. Kernel launch latency sets a high floor on kernel runtime
3. Relatively high cost of MPI getting data to/from GPU
4. Exposing sufficient parallelism to occupy GPU

Generic RV Performance – Weak Scaling

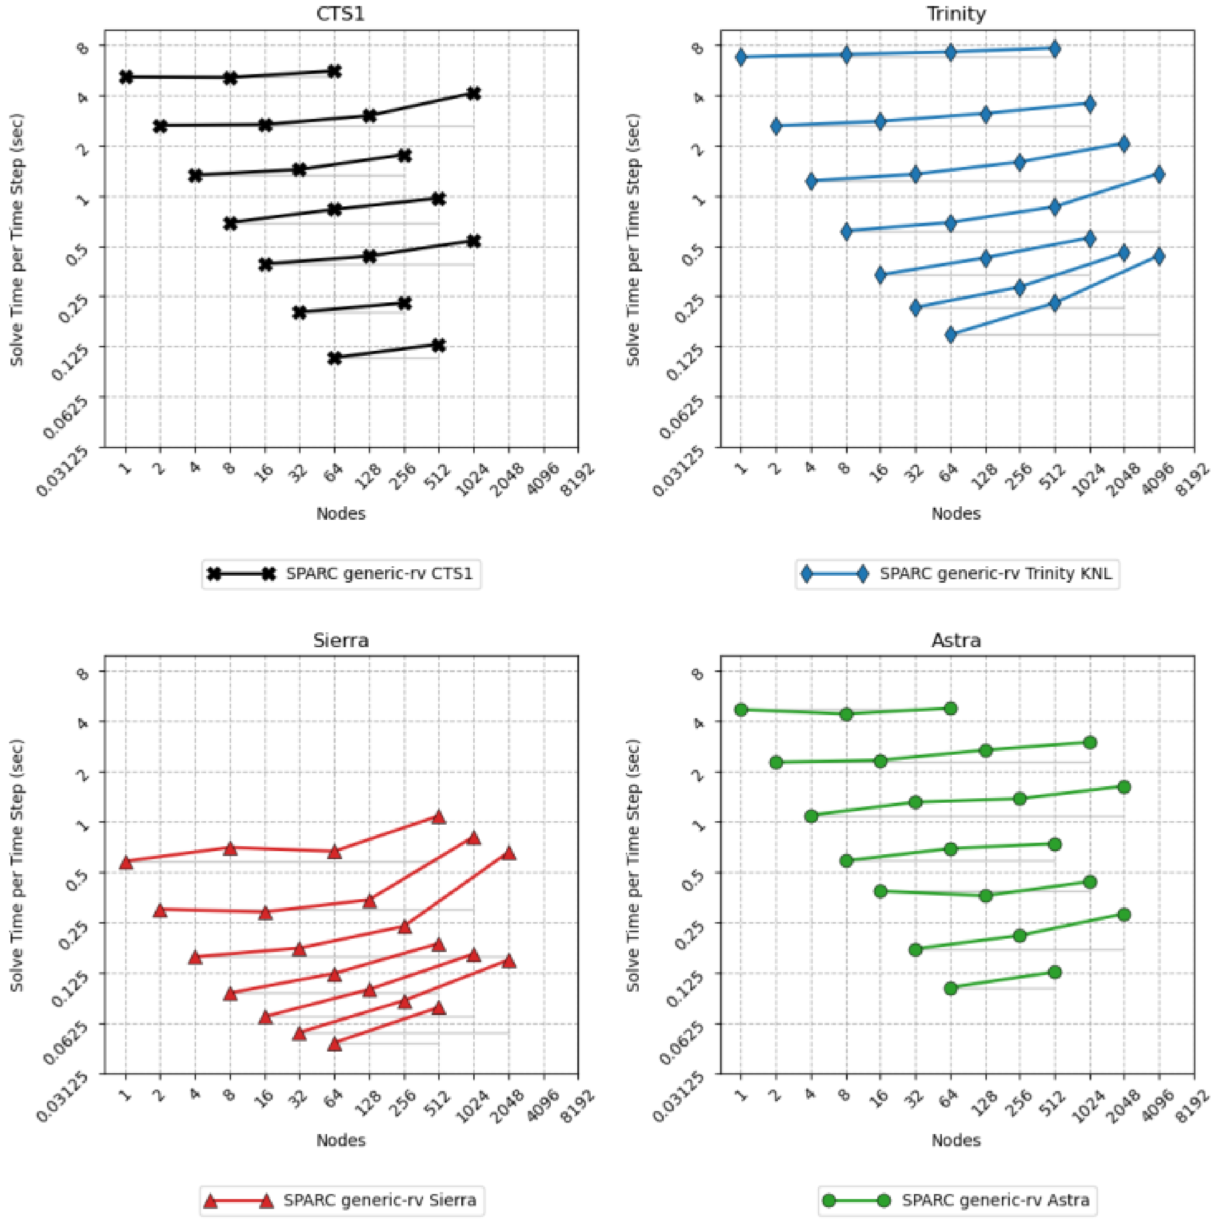

Trinity & Astra:
Good weak scaling with sufficient work per node

Sierra:
More work per node required for good weak scaling

Generic RV Performance – Throughput

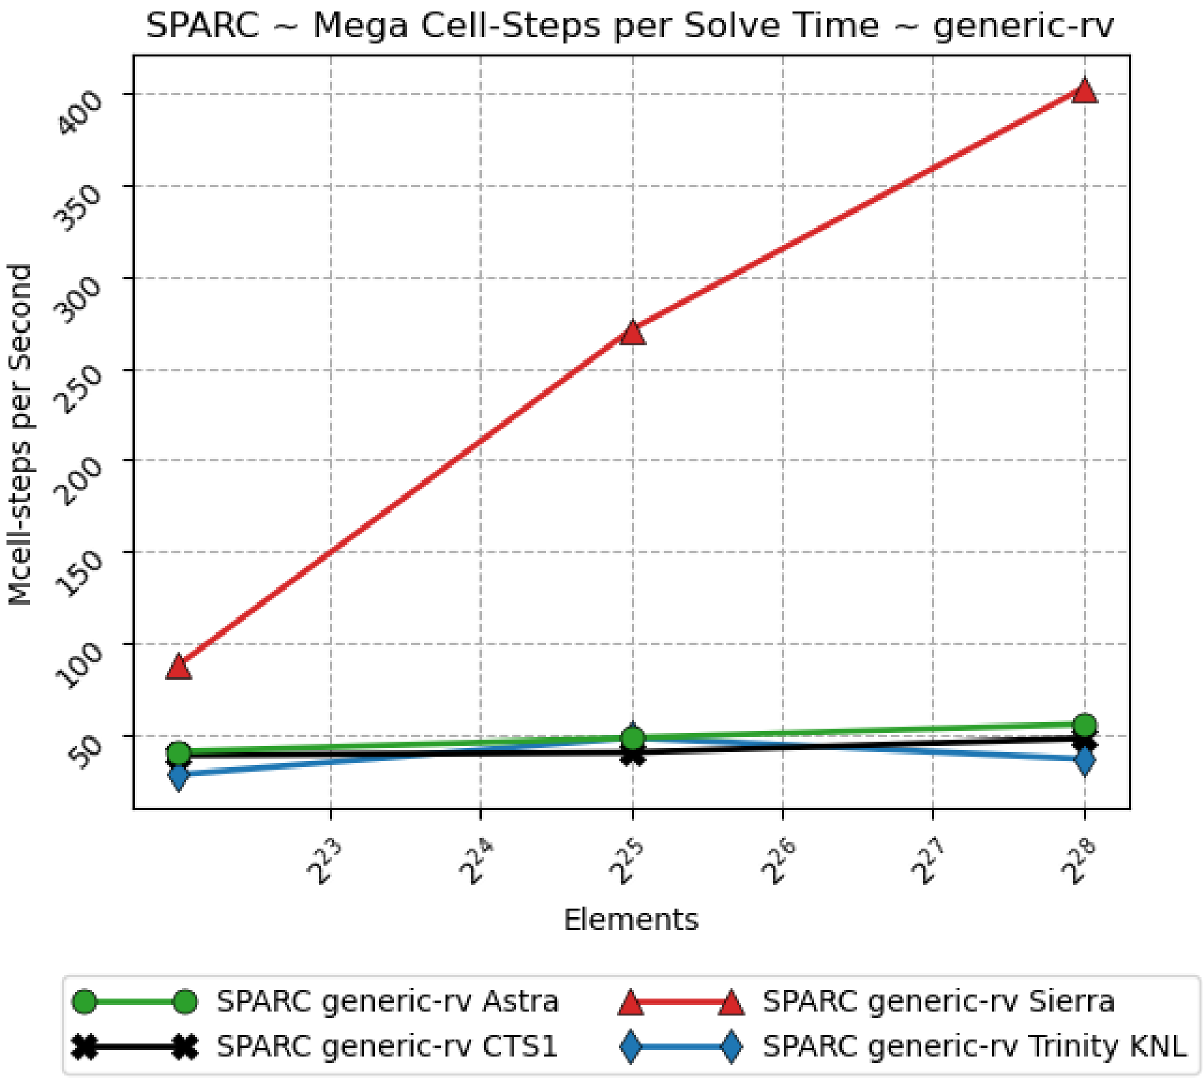

Throughput calculated from 64 node runs of R0, R1, R2 meshes.

Observations:

CTS1, Trinity/KNL, Astra/TX2:

- Close to independent of problem size

Sierra:

- GPUs thrive with more work

Summary

1. Successfully demonstrated SPARC scaling performance at scale on
   - 6144 nodes on Trinity
   - 2048 nodes on Astra
   - 2048 nodes on Sierra
2. SPARC achieves **excellent performance portability** and **speed-ups of up to**
   - 1.2x on Trinity
   - 1.4x on Astra
   - 8x on Sierra

   (1.2x on Trinity and 1.4x on Astra: when compared node-to-node with CTS-1)
3. SPARC achieves this with a **code base that is >95% platform agnostic**