SAND2015-10086C

Early Experiences with Node-Level Power Capping on the Cray XC40 Platform

E2SC Workshop @ SC15
November 15, 2015

Kevin Pedretti, Stephen Olivier, Kurt Ferreira
Center for Computing Research
Sandia National Laboratories

Galen Shipman
Los Alamos National Laboratory

Wei Shu
University of New Mexico

Sandia National Laboratories

Exceptional service in the national interest

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000. SAND NO. 2011-XXXXP

Outline

Sandia National Laboratories

- Background and Motivation
- Cray XC40 Power Management
- Experimental Approach
- Results
  - Don’t hit the power cap
  - Power usage can exceed the cap
- Next Steps and Conclusion

Large-Scale HPC is Power Intensive

- Trinity 2016 ATS-1 platform: 42 PF @ < 10 MW
  - Procurement contract: $174M
  - Five year energy bill: ~$50M
- APEX 2020 system targets:
  - Wall plate power < 20 MW
  - Peak power < 18 MW (max achievable in practice)
  - Nominal power < 15 MW (for APEX workloads)
- Multi-megawatt transient swings a possible issue
- Current practice mostly ignores power management
  - Run system at max performance (max power) settings all the time
  - No real incentive to actively manage power, other than for research
- Expect future large-scale systems to be power constrained

APEX

Alliance for Application Performance at Extreme Scale

LANL, LBL, SNL

Power budget is a critical resource
Need to direct it to where it has most benefit

Many Power Management Techniques

Sandia National Laboratories

- P-states, C-states, T-states, turbo boost, slack time throttling, power caps, concurrency throttling, GPU/CPU power shifting, hardware overprovisioning, mem throttling, net throttling, …

- Key questions:
  - Which techniques are effective for large-scale HPC systems?
  - Can we actually deploy these on large-scale systems + show benefit?
  - How do we get vendors, system operators, and users on board?

- Vendors starting to provide system-level power measurement and control capabilities; how well do they work? ☺

We evaluate Cray's power capping infrastructure, intended for use by workload managers and admins

Power Capping

- Software specifies desired power limit, hardware enforces
- Hardware continuously monitors power draw, takes action if power usage over limit (e.g., lower CPU freq, clock gating)
- Usually implemented as avg. power over sliding time window
- Initial work in HPC context by Rountree'12 et al. using RAPL

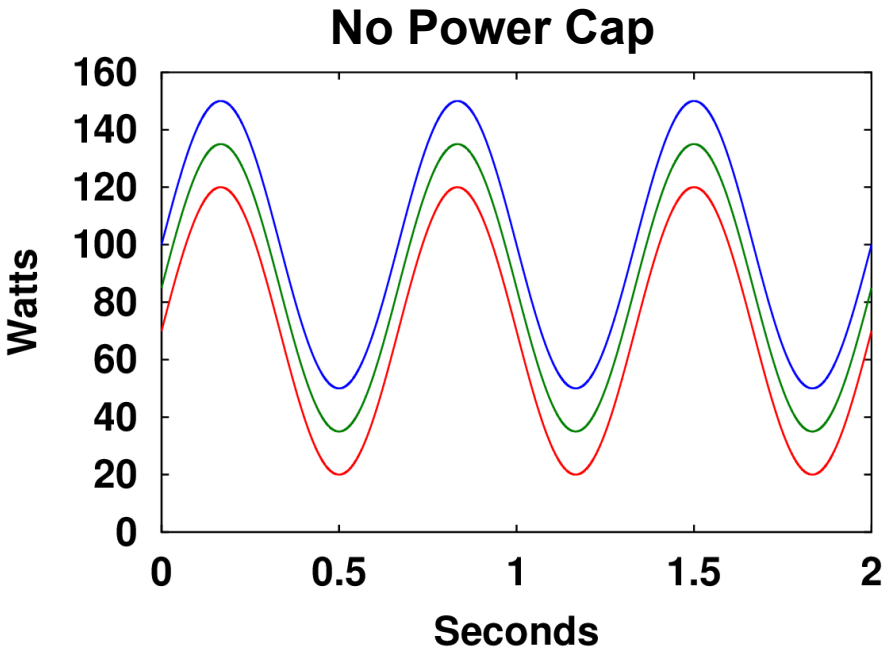

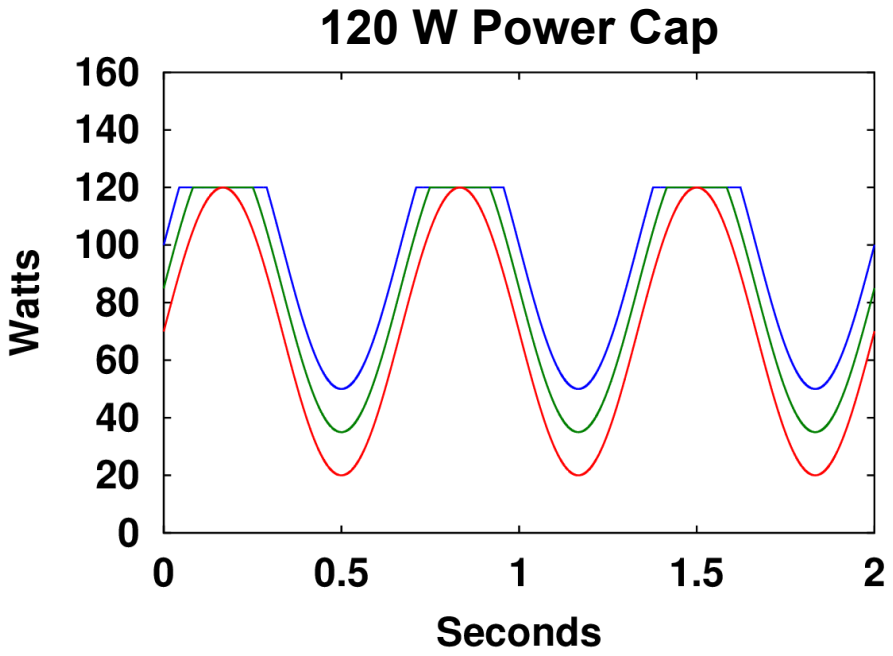

Power Caps Provide Predictable Power

- Power capping appealing due to max power draw "guarantee"
  - Recapture (*WallPlate – Pcap*) headroom and reallocate elsewhere
  - Potentially useful for doing power aware scheduling
- P-state control appealing due to predictable performance, power unknown

Single-node HPL Variability Across 100 Nodes

P-State Sweep (Fixed Performance, Variable Power)

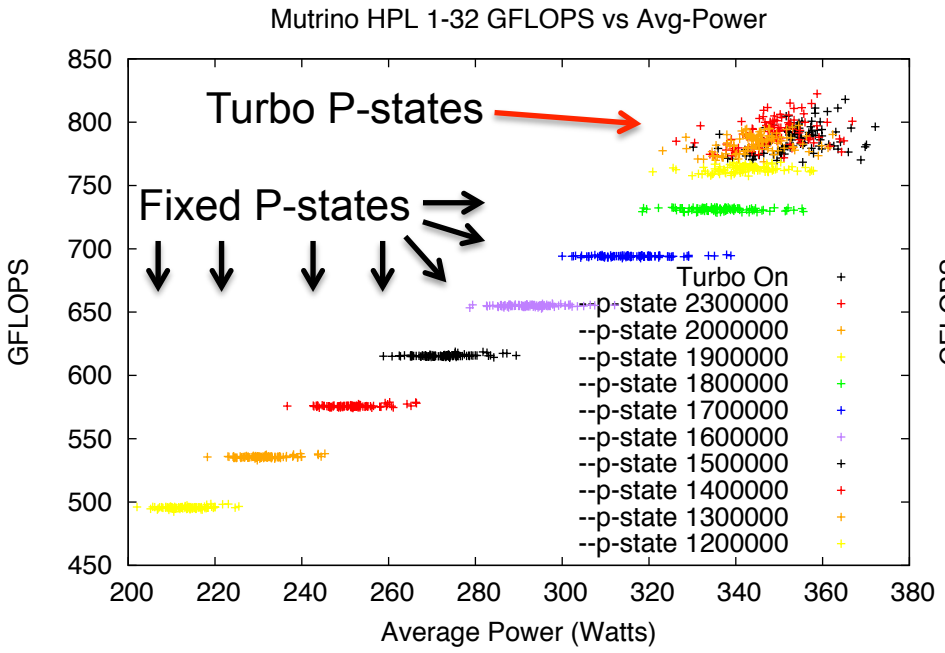

Power Cap Sweep (Fixed Power, Variable Performance)

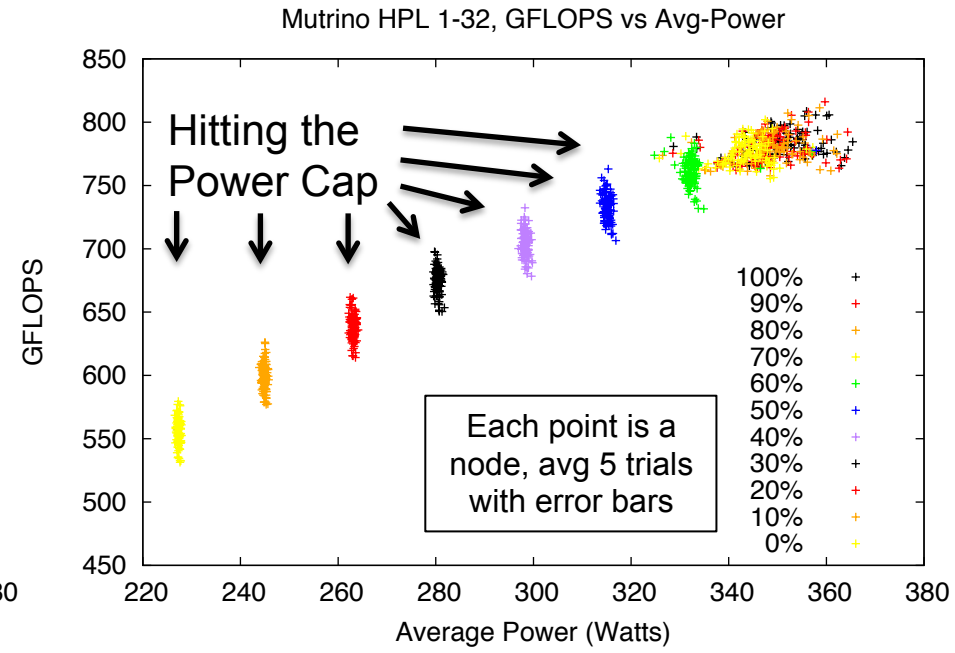

Outline

Sandia National Laboratories

- Background and Motivation
- Cray XC40 Power Management
- Experimental Approach
- Results
  - Don’t hit the power cap
  - Power usage can exceed the cap
- Next Steps and Conclusion

Node-Level Power Capping

Sandia National Laboratories

Trinity "Haswell" Compute Node

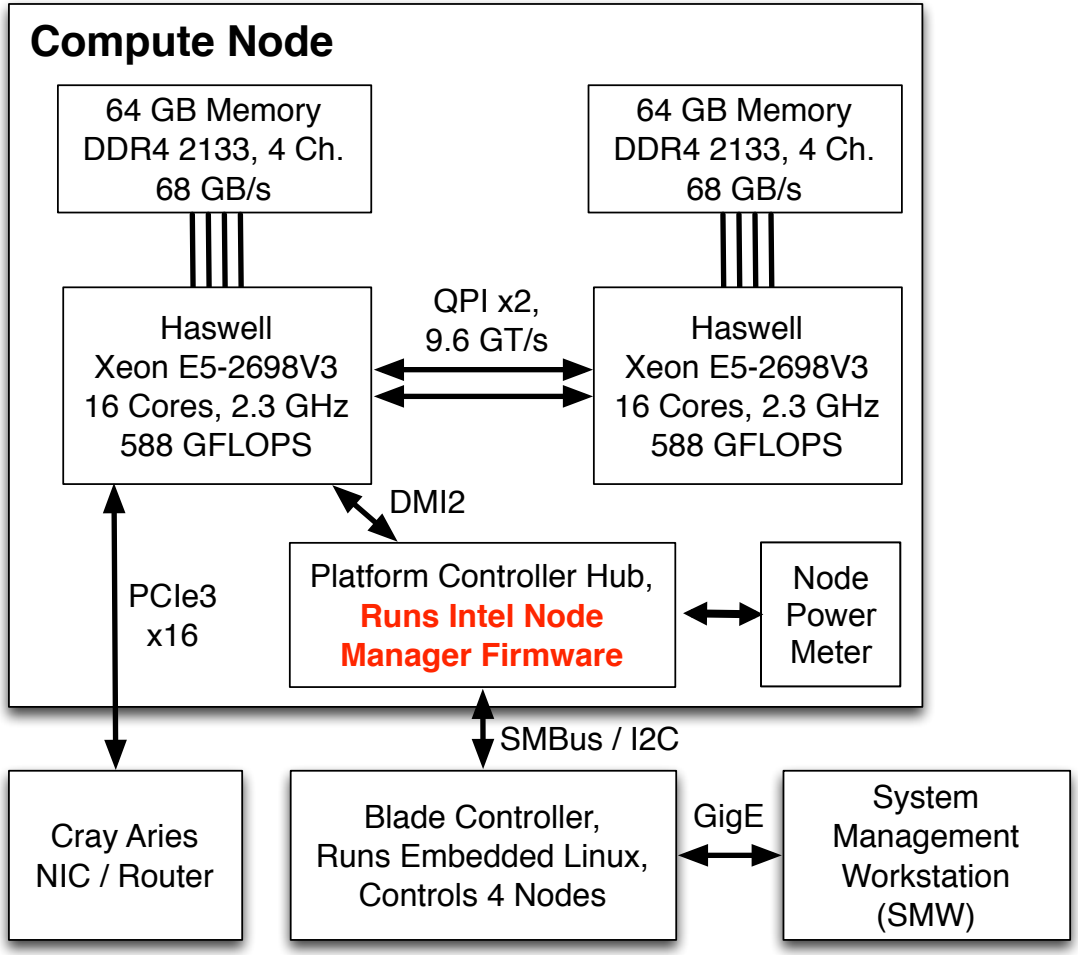

- Blade controller receives power cap command from management network
- Blade controller relays command to Intel PCH
- PCH runs Intel's Node Manager firmware, which enforces the power cap
- Cray node-level power meter provides feedback to Node Manager
- Node manager makes use of each socket's RAPL
- Power capping algorithm not disclosed by Intel

Specification: Intel Intelligent Power Node Manager 3.0, March 2015, Doc# 32200-001US

Job/System-Level Power Capping

Sandia National Laboratories

Cray XC40 Power Management Architecture

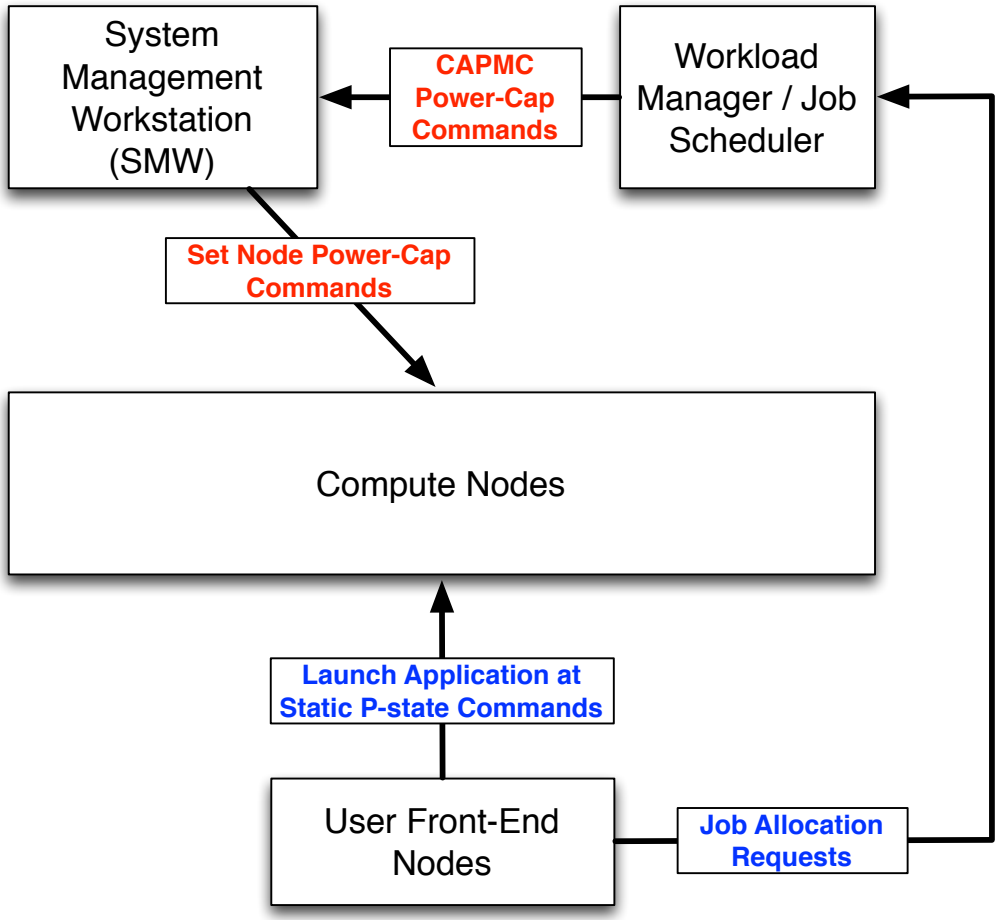

- A single management workstation controls system, the SMW
- Node-level power caps set from SMW, distributed to compute nodes via out-of-band management network
- Admins use `xtpmaction` command to set power caps manually
- Workload managers use Cray's CAPMC web API to set power caps
- Users may launch their job at a fixed p-state, default is P0 (turbo on)

See: Monitoring and Managing Power Consumption on the Cray XC System, April 2015, S-0043-7203

Power Management Limitations

Sandia National Laboratories

- Setting power cap requires administrative privilege
- Users cannot dynamically change P-states
- Users cannot dynamically change Power Caps
- Users cannot access RAPL, other than via read-only PAPI
- Users cannot access Cray Power Management Database
- 1 Hz power samples for each node collected by default, configurable up to 5 Hz for small sets of nodes

Outline

Sandia National Laboratories

- Background and Motivation
- Cray XC40 Power Management
- Experimental Approach
- Results
  - Don’t hit the cap
  - Power usage can exceed the cap
- Next Steps and Conclusion

Testing Procedure

Sandia National Laboratories

- Test system was "Mutrino" Cray XC40 @ Sandia, Mini Trinity
  - 100 Haswell compute nodes, identical to Trinity Haswell nodes
  - Testing performed in dedicated mode with no other users
  - Always used 32 MPI processes per node, 1-1 pinned to physical CPUs
- Test configurations cross-product of:
  - 4 workloads (HPL, HPCG, S3D, CTH)
  - 5 scales (1, 8, 32, 64, 96 nodes)
  - 8 p-states (1.2, 1.4, 1.6, 1.8, 1.9 2.0, 2.3, 2.301 GHz)
  - 5 power caps (230, 276, 322, 369, 415 watts)
  - 3 trials each configuration

> 2400 Runs

- For each trial, collected
  - Energy consumed and total runtime from Cray RUR
    (energy consumed includes CPUs, memory, PCH; excludes Aries NIC)
  - 10 Hz power samples and CPU 0 freq samples using Power API

powerapi.sandia.gov

Workloads

Sandia National Laboratories

- Weak scaling (scaled problem size with node count)
- Top 500 Benchmarks
  - HPL – High Performance Linpack
    - Compute intensive
    - 24 GB per node (N = 549K for 96 nodes)
  - HPCG – High Performance Conjugate Gradient
    - Memory intensive
    - 30 GB per node (104x104x104)
- Application Benchmarks
  - S3D – Combustion Simulation
    - Network intensive, 3-D nearest neighbor exchanges
    - 48^3 gridpoints per node
  - CTH – Shock Physics Hydro Code
    - Memory and network bandwidth intensive
    - Shaped charge input deck, ~ 50 GB per node

“Mutrino” Test System

- Cray XC40 with 100 compute notes
- Trinity lookalike system, but smaller
- Power cap range 230 to 415 W
- P-state range 1.2 to 2.3 GHz
- Max turbo 3.6 GHz, AVX freq 1.9 GHz

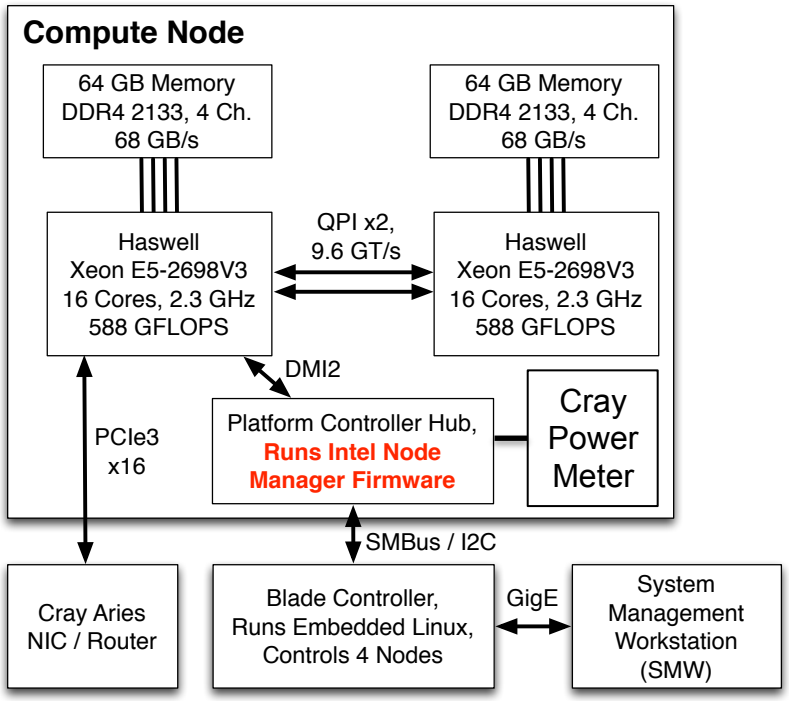

Power Cap Settings Tested (5)

Cray Power Cap Setting	Power Cap Per-Node (Watts)	Savings Potential (Watts)	Savings Potential (Percentage)
No Cap	~ 415 W	N/A	N/A
75%	369 W	46 W	11%
50%	322 W	93 W	22%
25%	276 W	139 W	33%
0%	230 W	185 W	45%

P-State Settings Tested (8)

Cray P-state Name	Clock Frequency (GHz)	Percent of Peak
2301000	2.3–3.6 (Turbo On)	> 100%
2300000	2.3	100%
2000000	2.0	87%
1900000	1.9	83%
1800000	1.8	78%
1600000	1.6	70%
1400000	1.4	61%
1200000	1.2	52%

Outline

Sandia National Laboratories

- Background and Motivation
- Cray XC40 Power Management
- Experimental Approach
- Results
  - Don’t hit the power cap
  - Power usage can exceed the cap
- Next Steps and Conclusion

S3D No Cap: "Things Look Good"

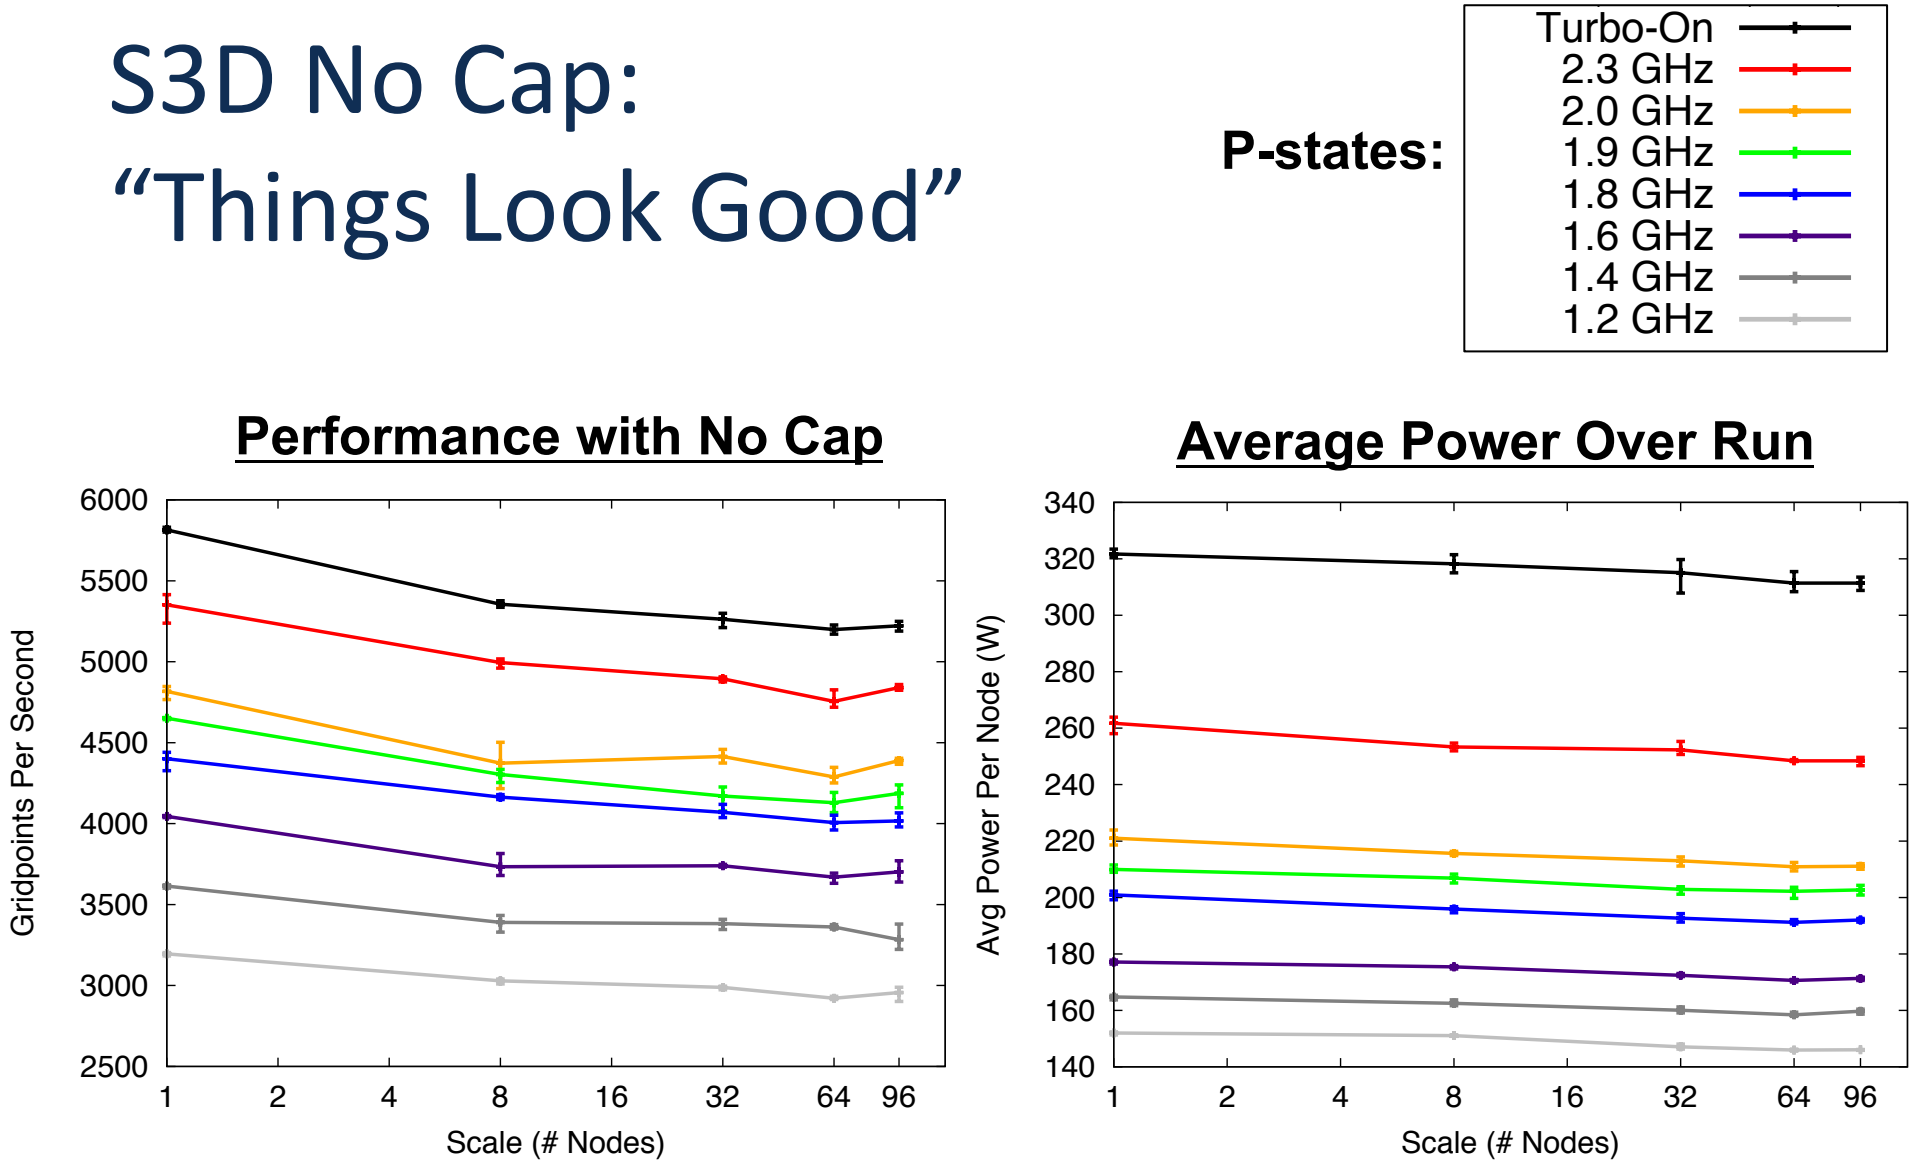

S3D 322W Cap (50%): "Turbo Hits the Cap"

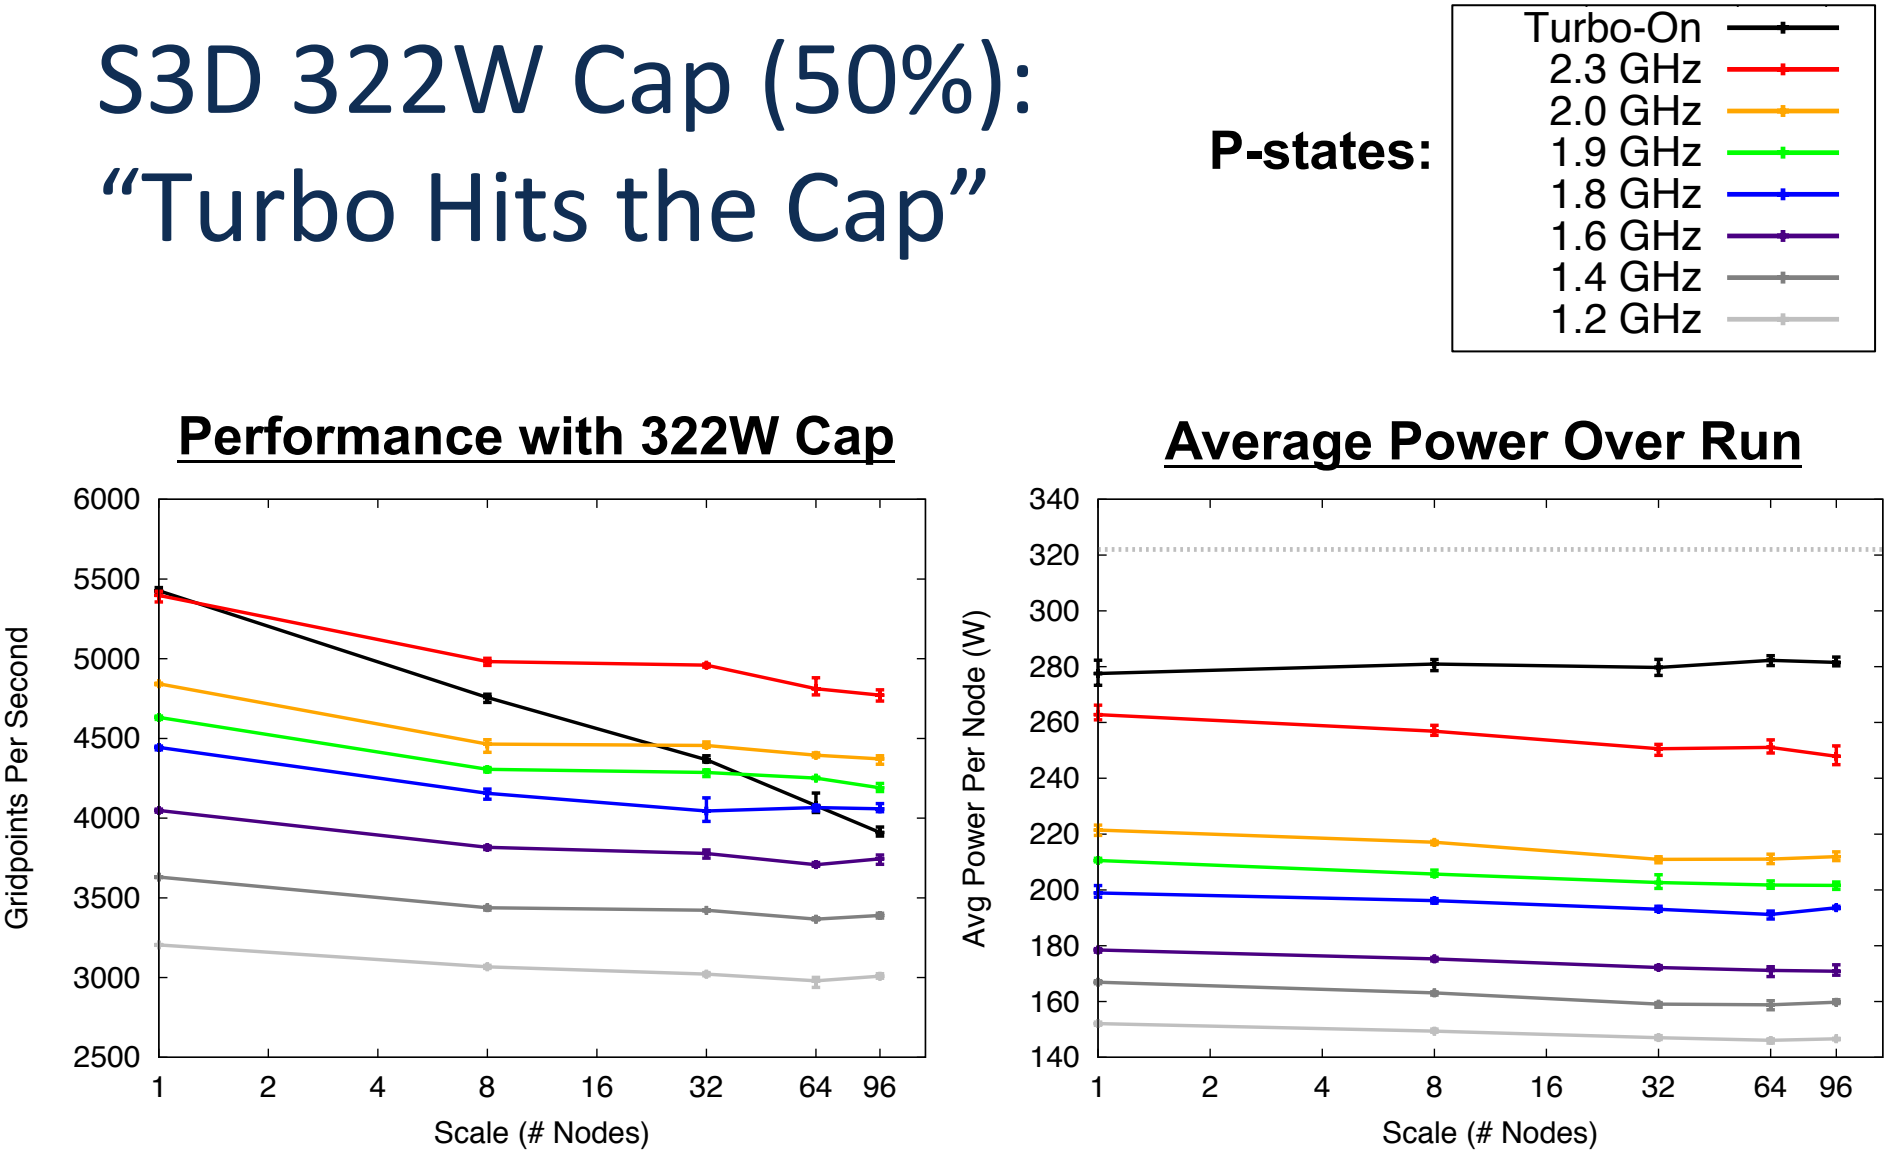

S3D 276W Cap (25%): "2.3 GHz Hits the Cap"

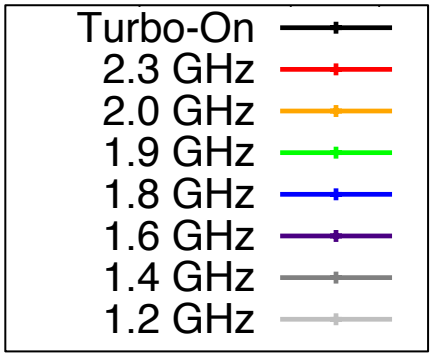

Performance with 276W Cap

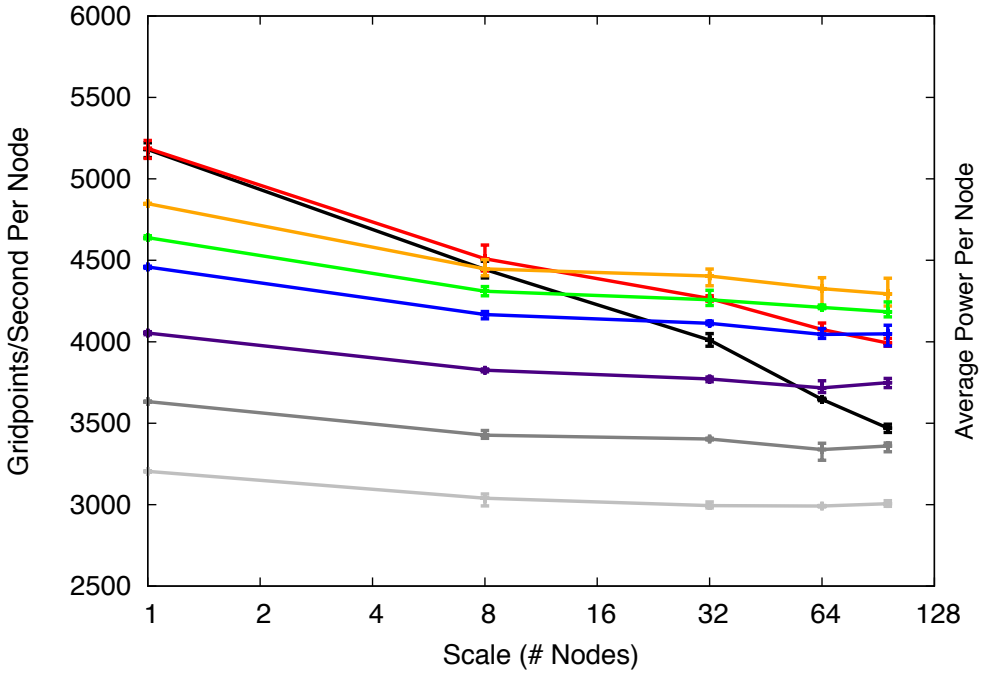

Average Power Over Run

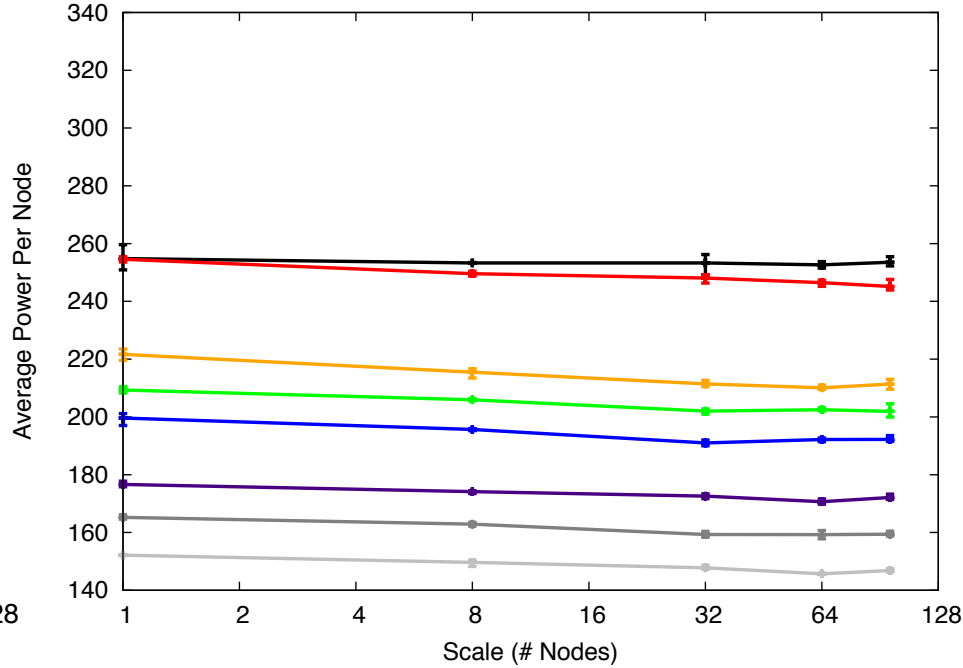

S3D 230W Cap (0%): "> 1.9 GHz Hits the Cap"

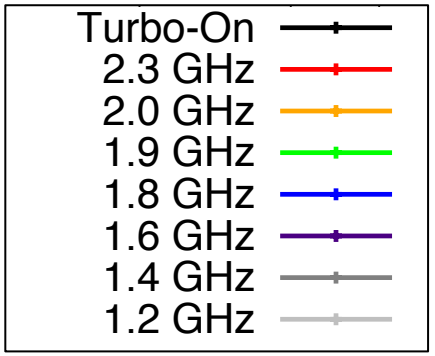

Performance with 230W Cap

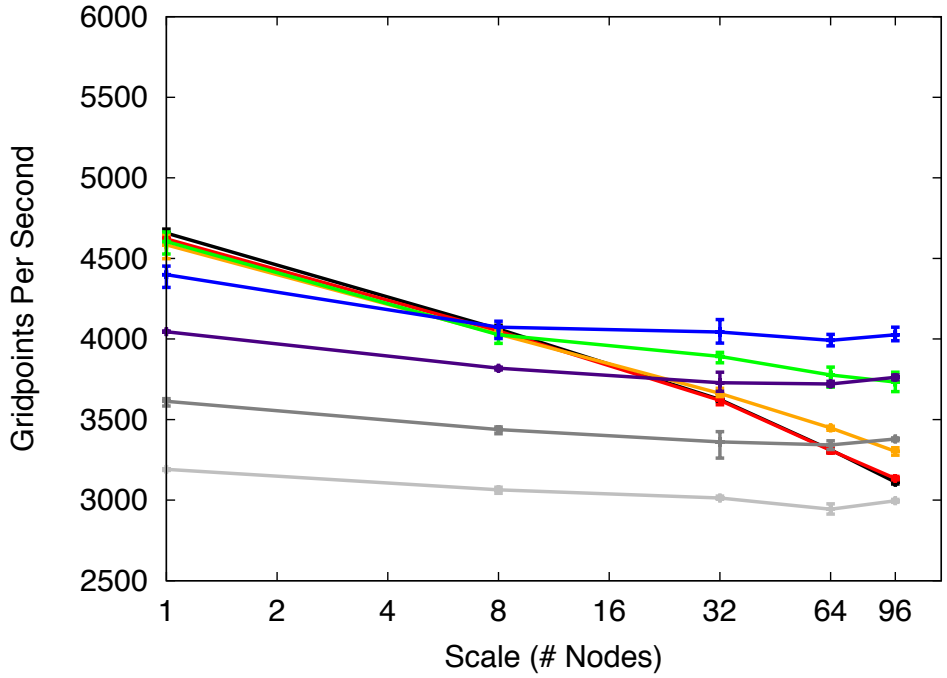

Average Power Over Run

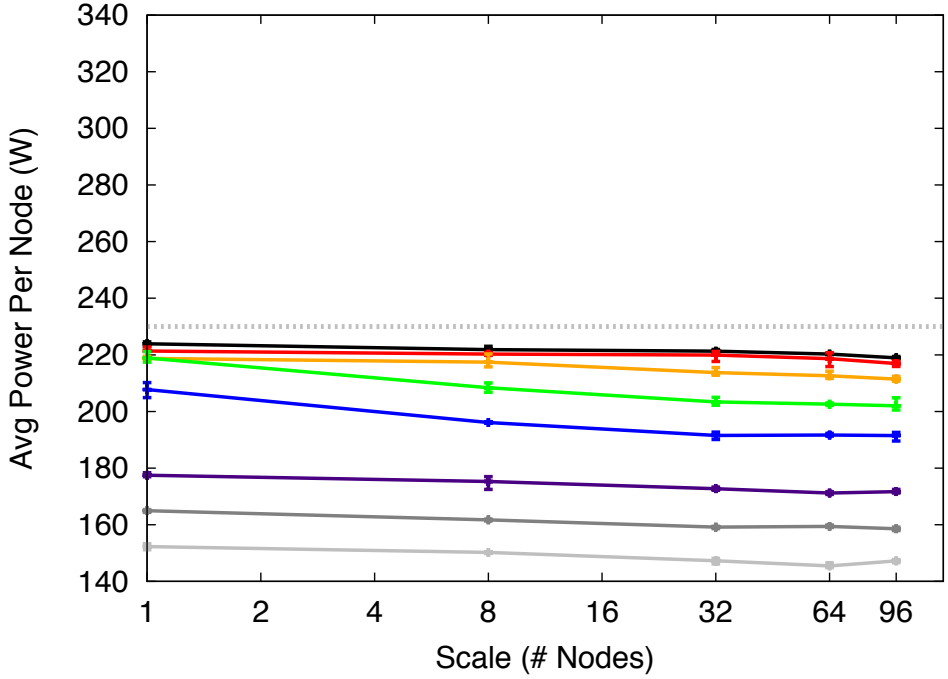

Hitting cap leads to increasing performance degradation with scale

S3D Summary

Sandia National Laboratories

- As power cap increases, the highest N p-states experience performance degradation with scale
- This seems to make sense:
  When you hit the cap, thus triggering power throttling, bad things start to happen
- What about the other apps?

Sandia National Laboratories

CTH

CTH No Cap: "Things Look Good"

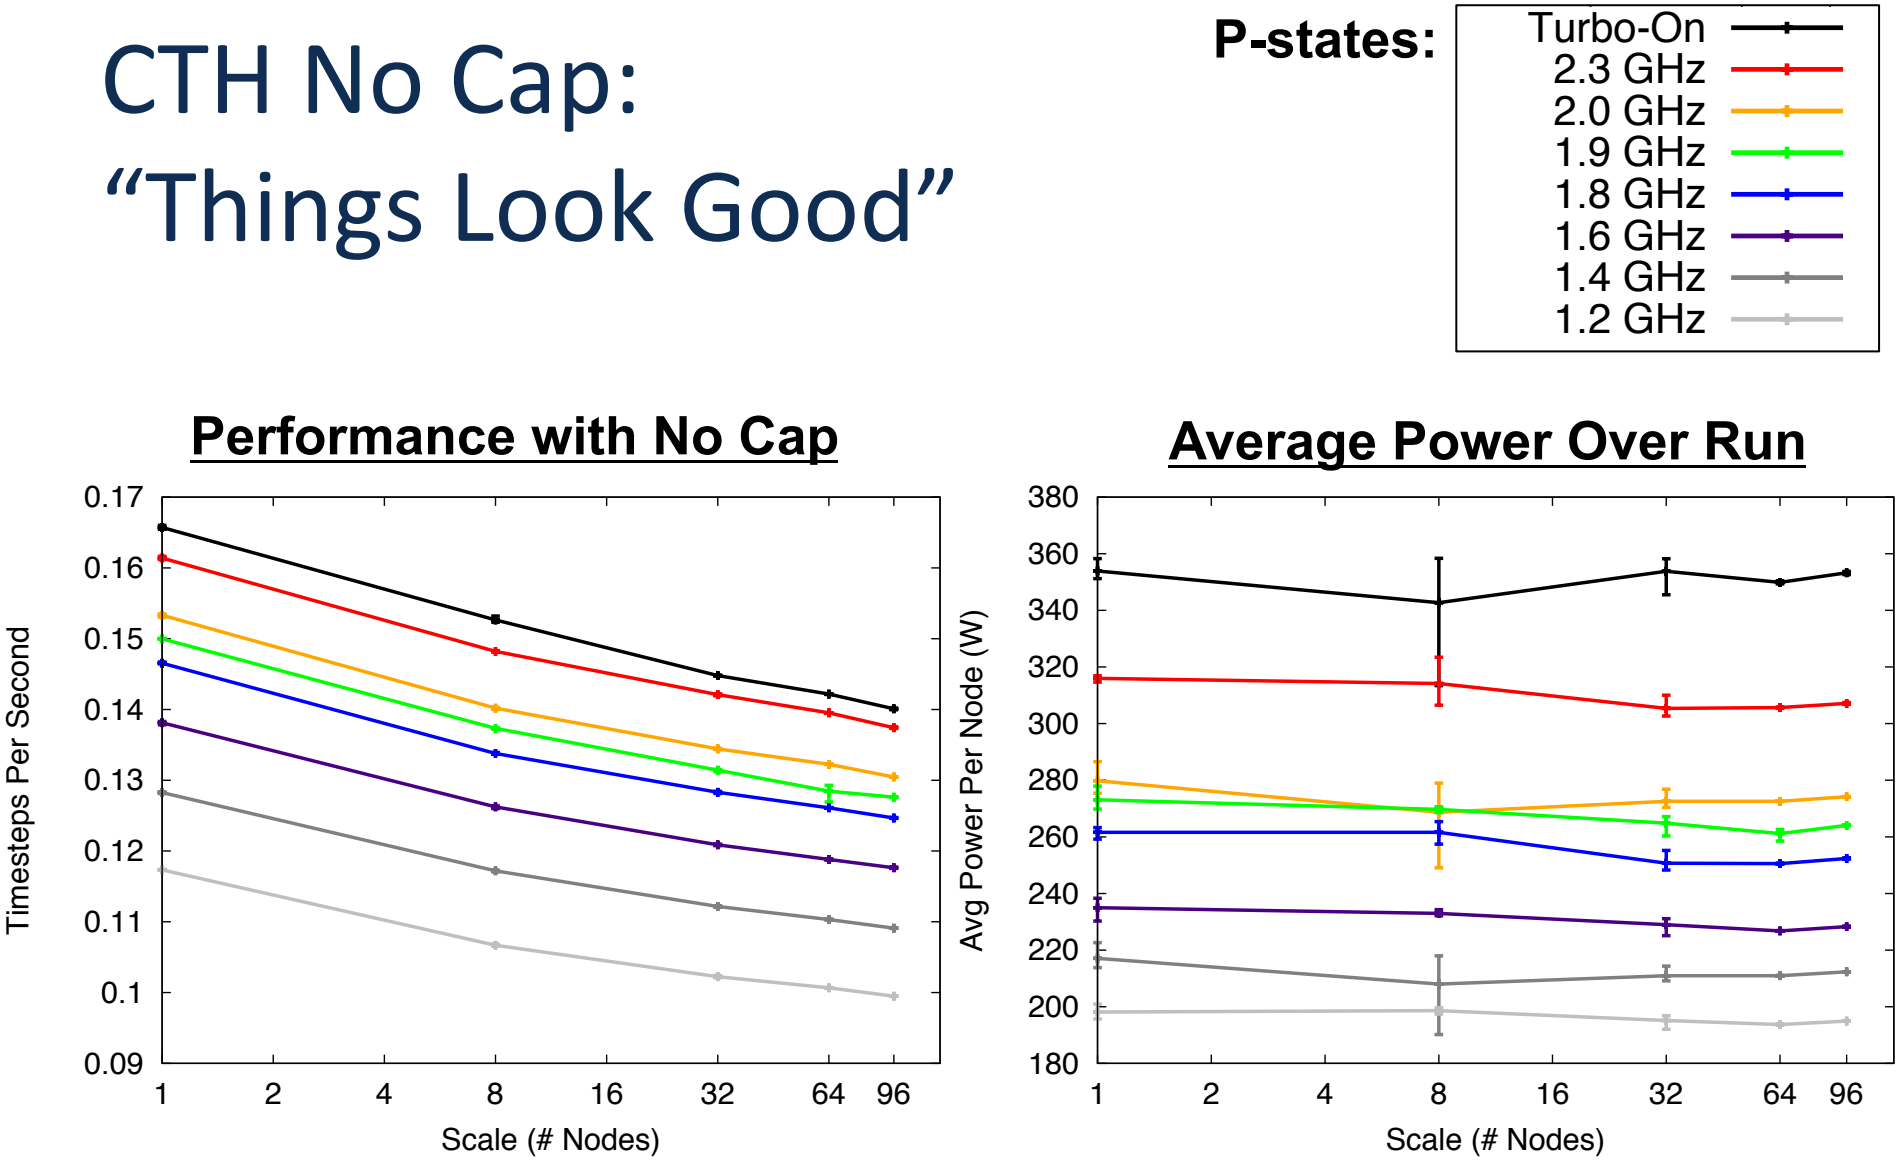

CTH 322W Cap (50%): "Turbo & 2.3 GHz Hit Cap"

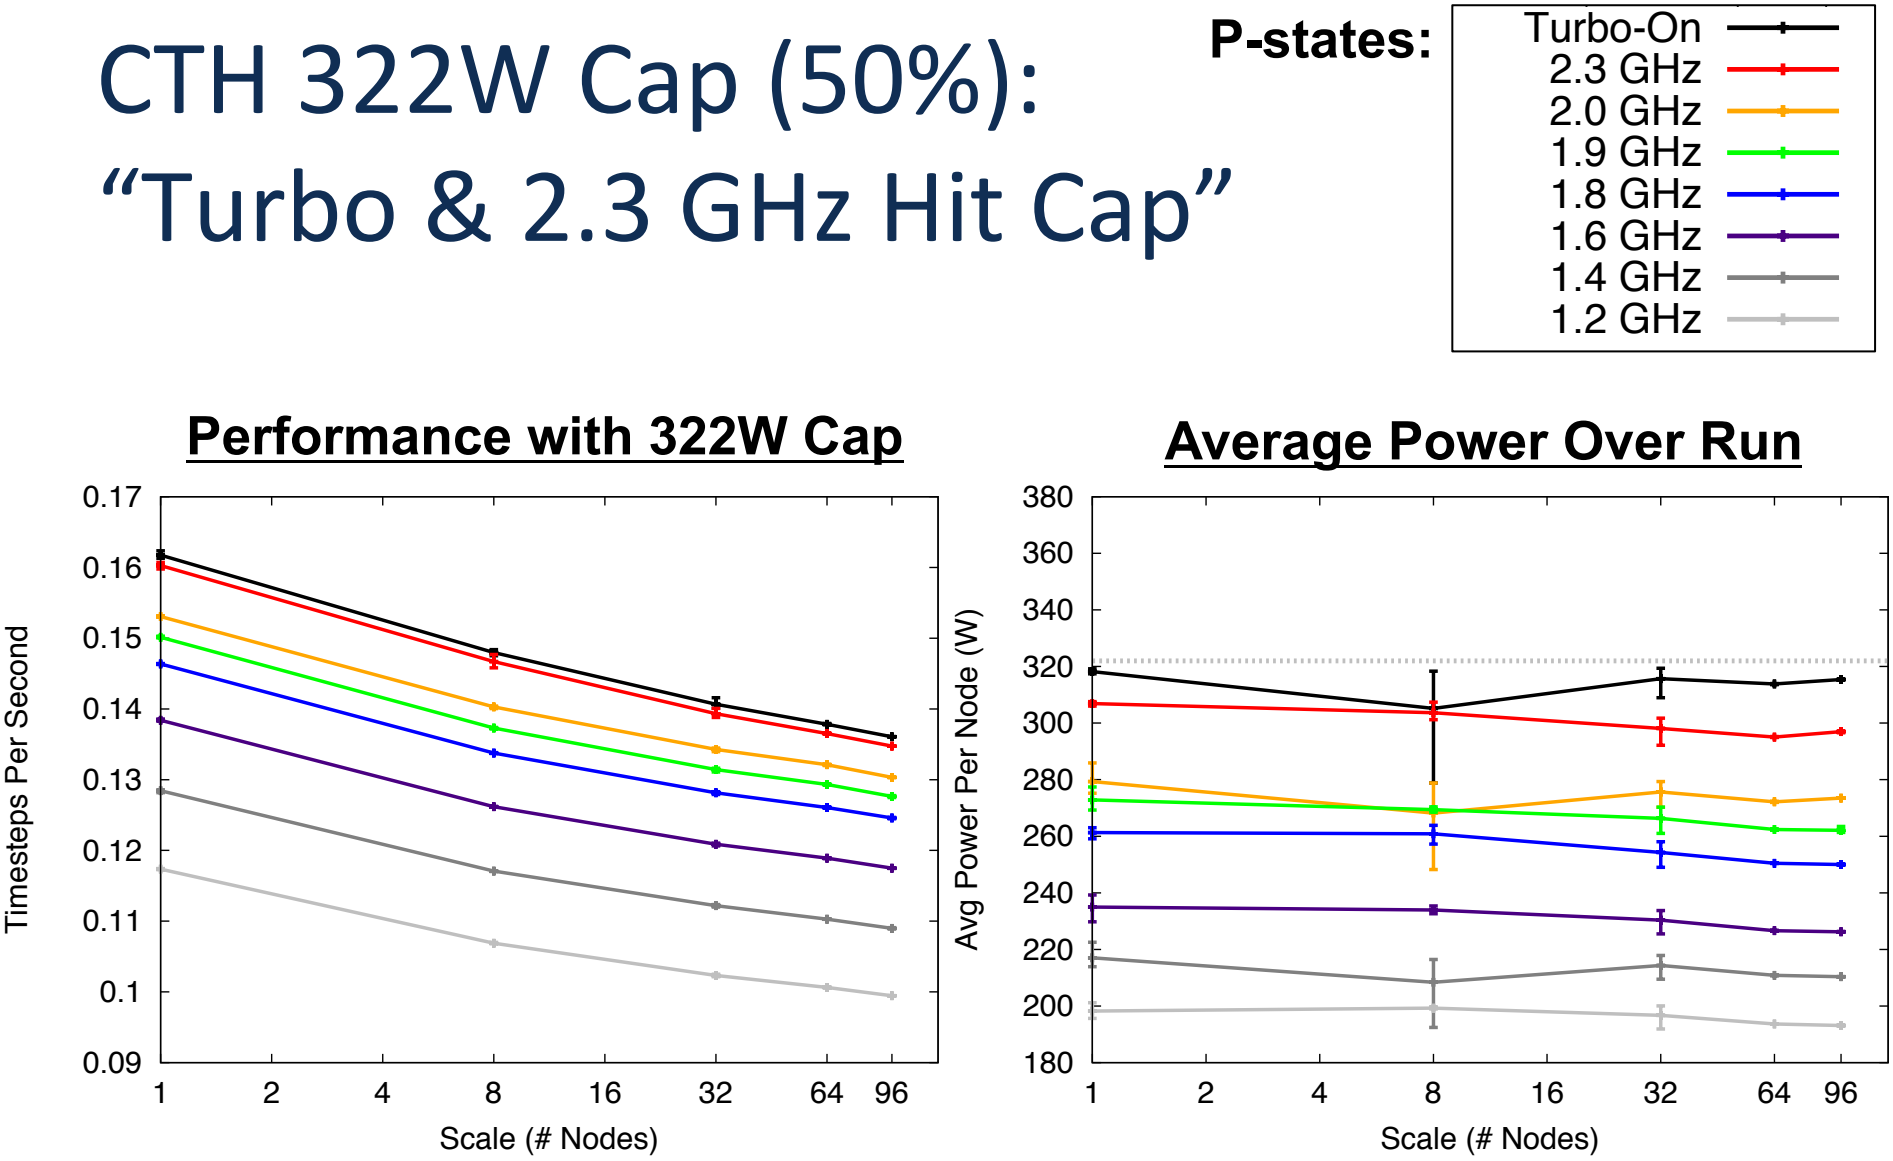

CTH 230W Cap (0%): "> 1.6 GHz Hit Cap"

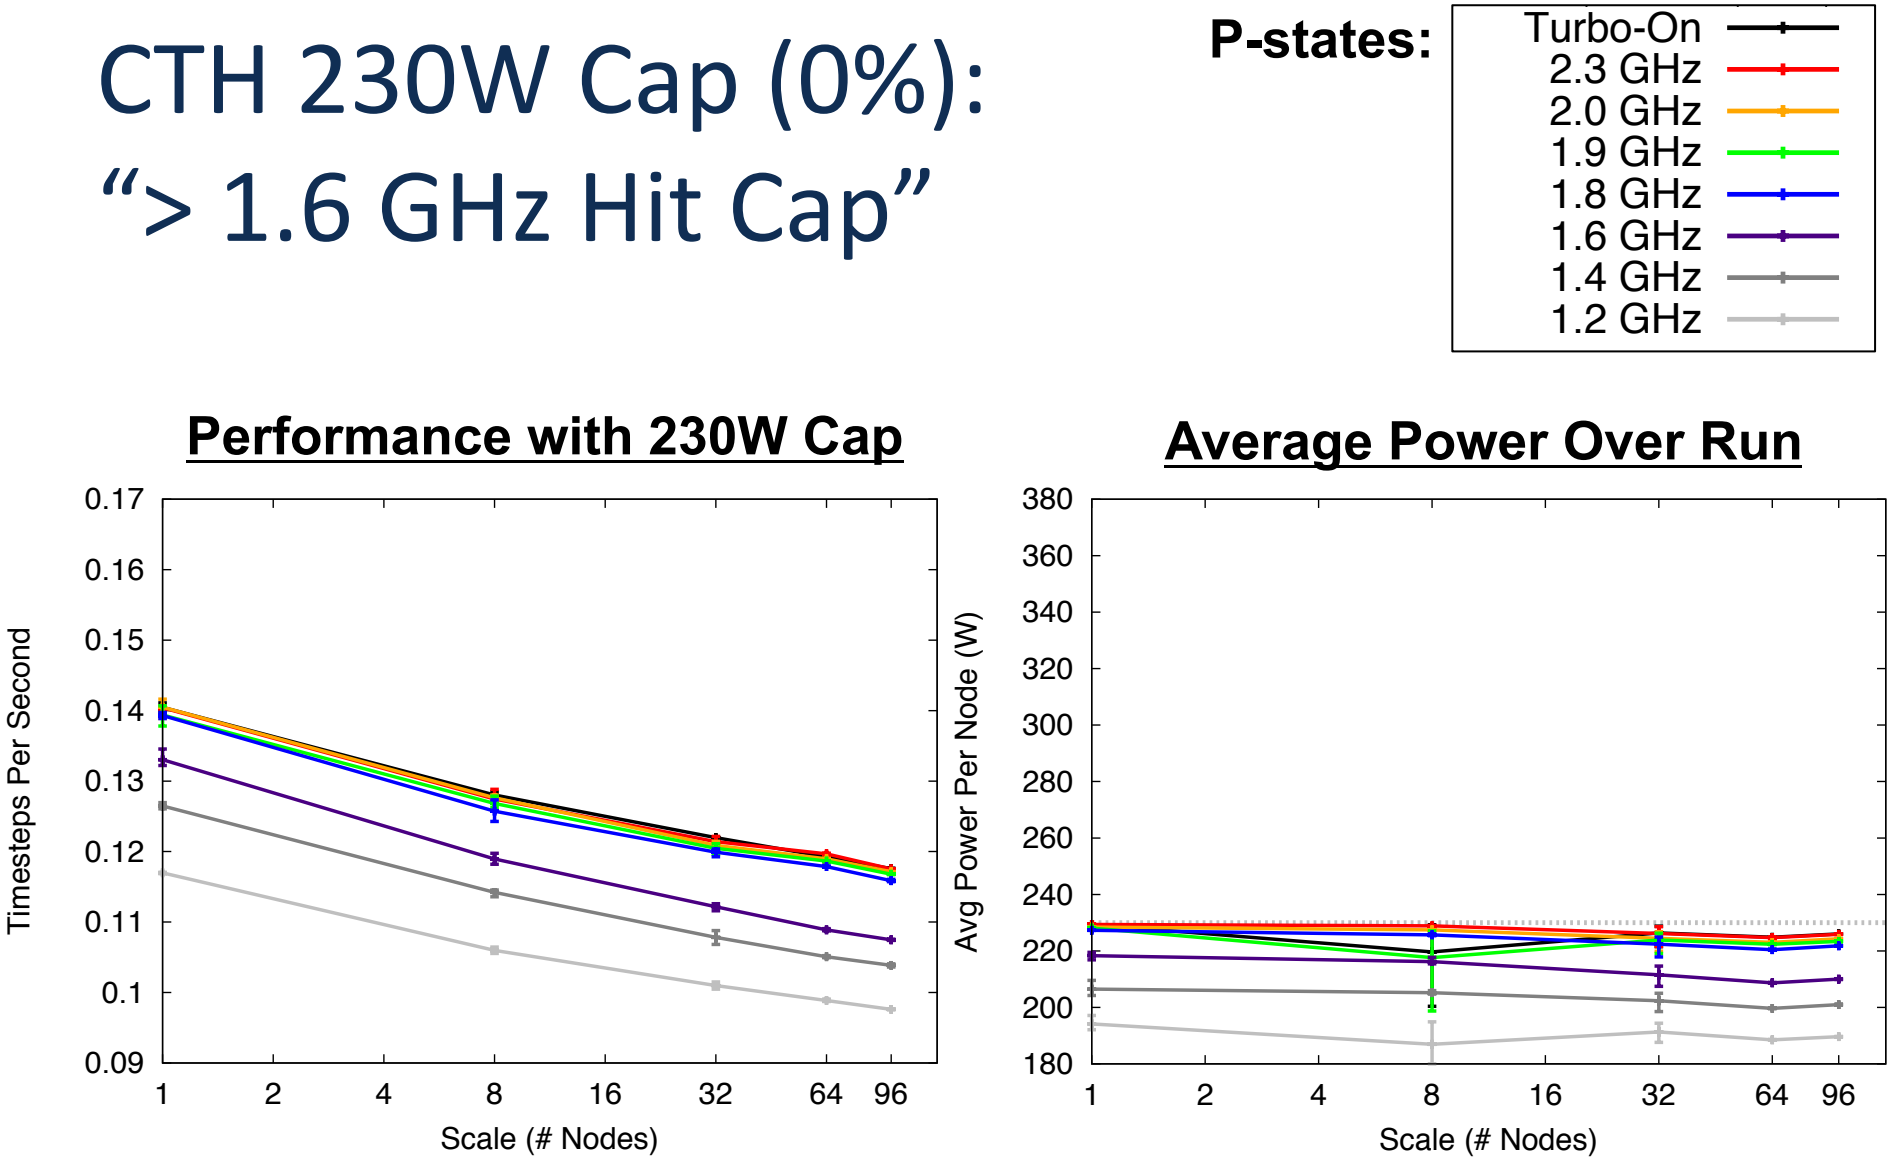

Unlike S3D, CTH does not experience performance degradation with scale when it hits the cap

Sandia National Laboratories

HPL

HPL No Cap: "Things Look Good"

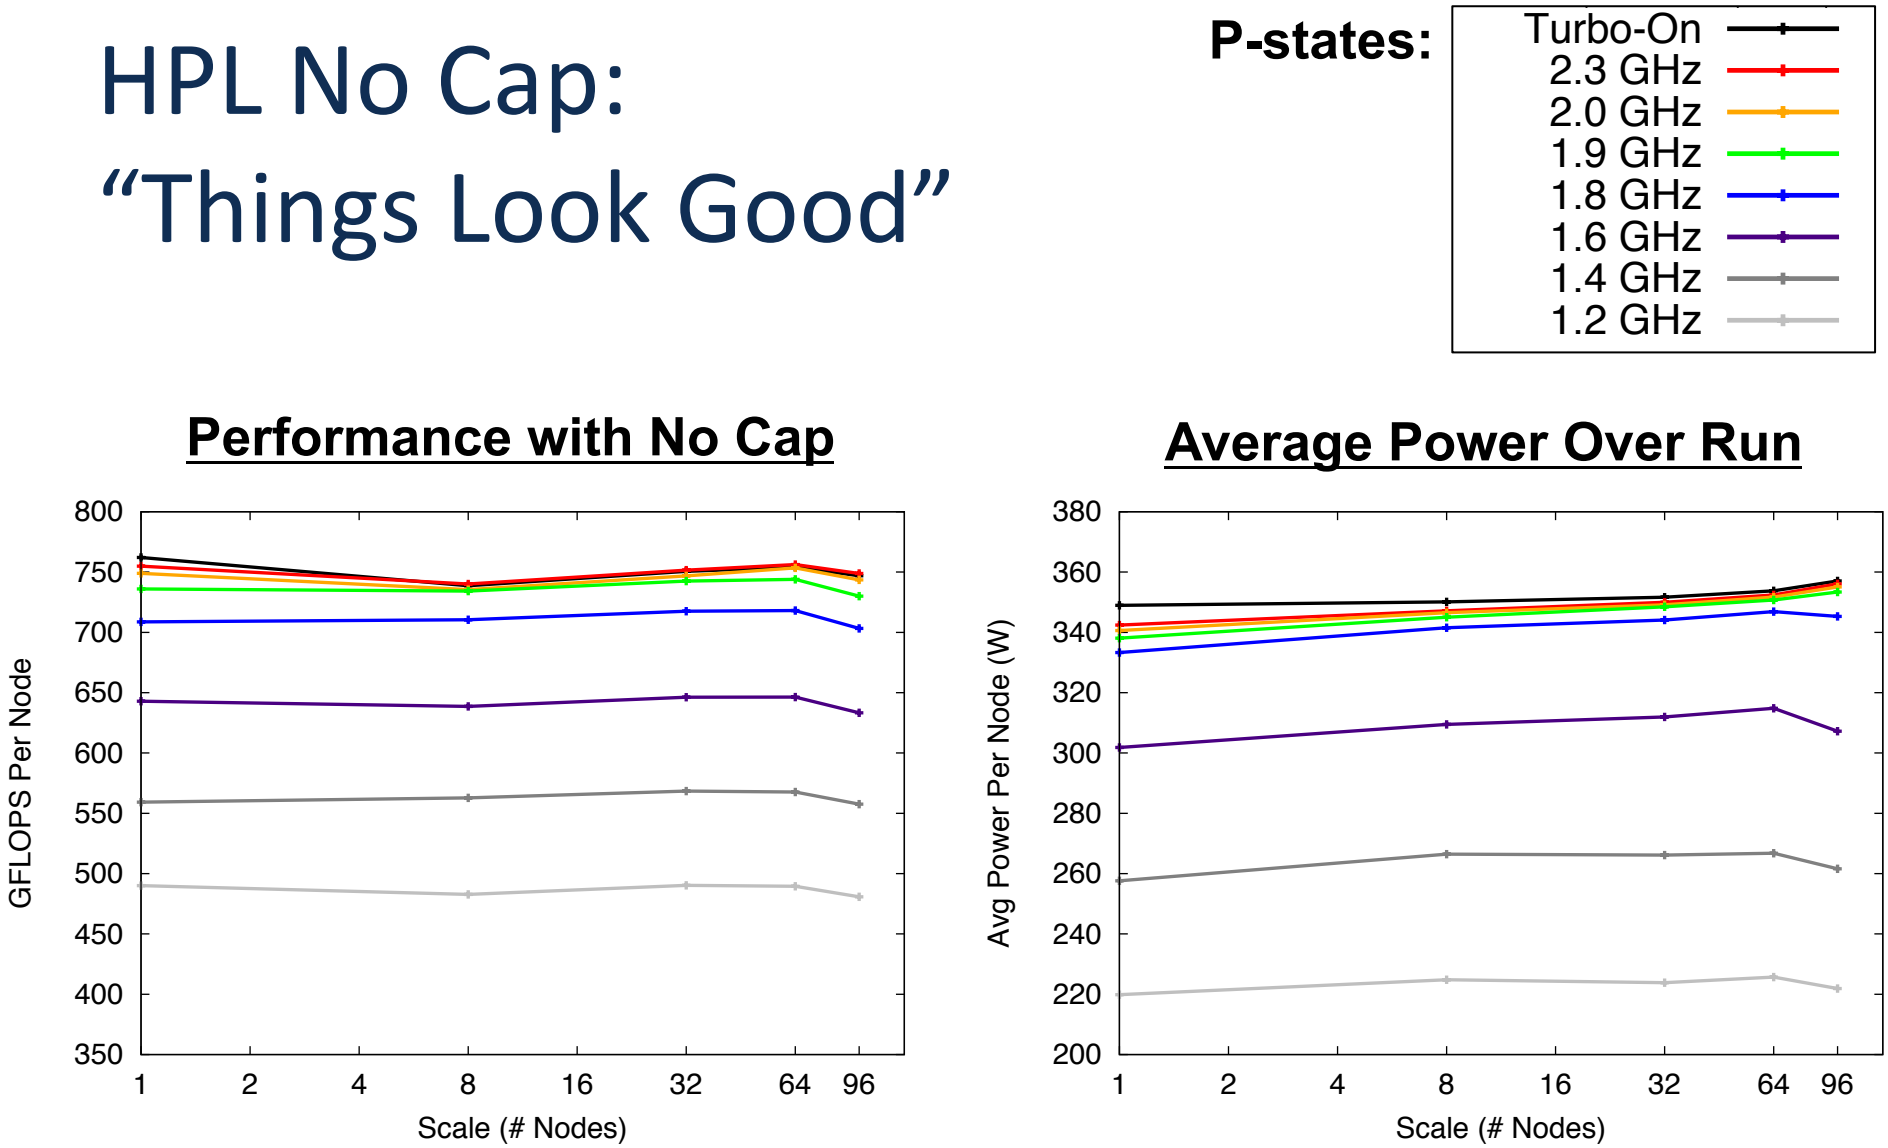

HPL 322W Cap (50%): "> 1.6 GHz Hit Cap"

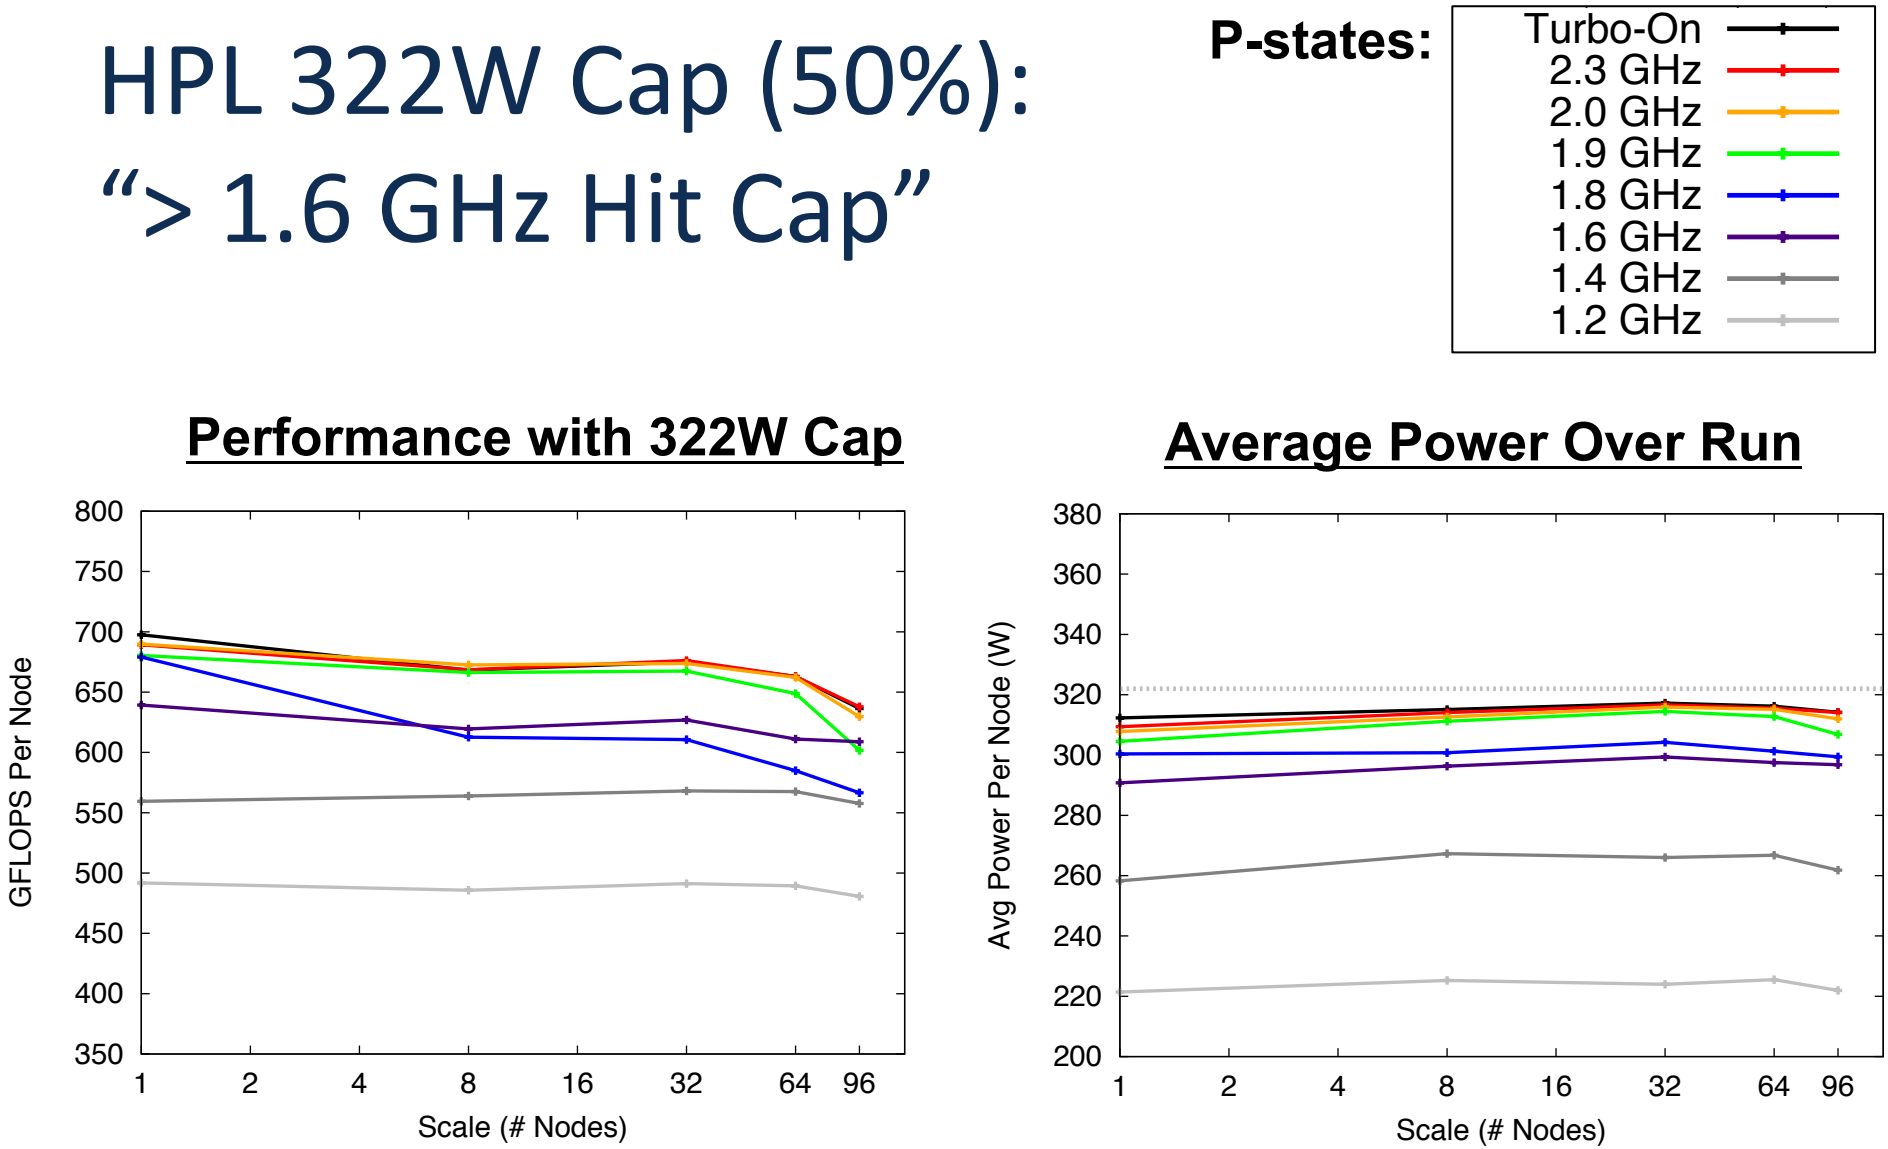

HPL 230W Cap (0%): "All P-states Hit Cap"

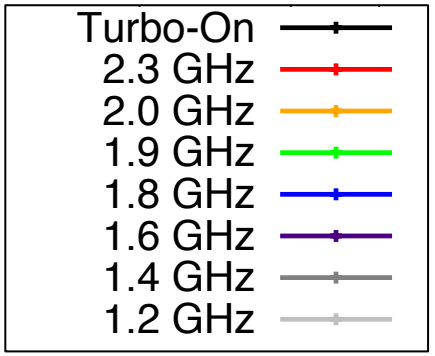

Performance with 230W Cap

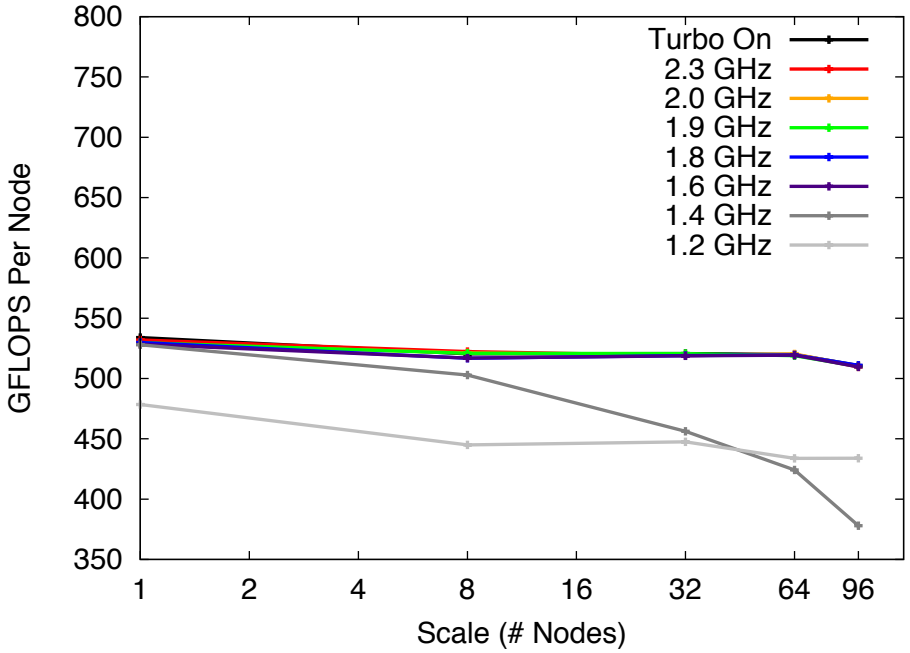

Average Power Over Run

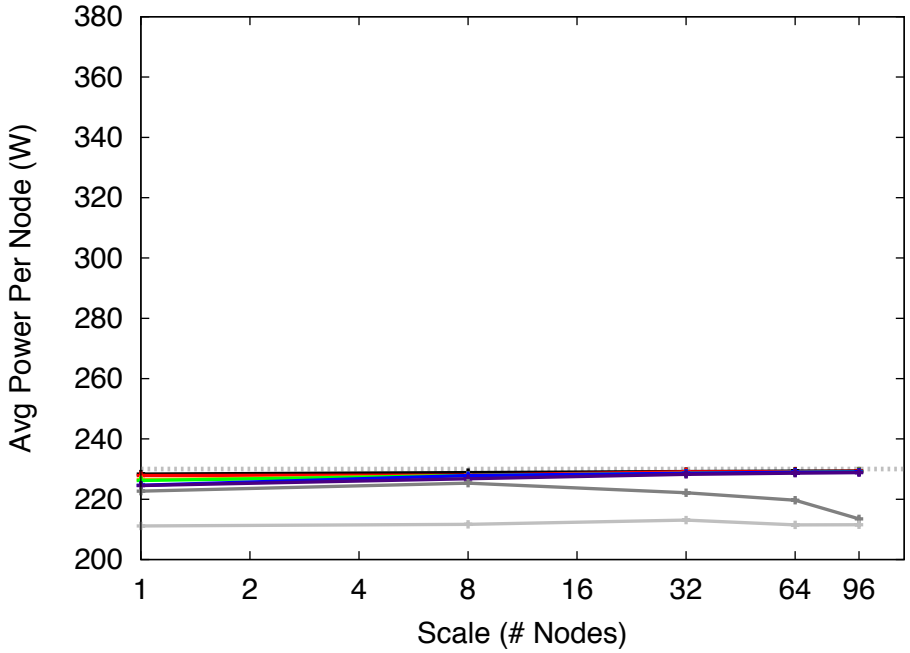

1.4 GHz only experiences performance degradation with scale, performance drops below 1.2 GHz floor!

Sandia National Laboratories

HPCG

HPCG No Cap: "Things Look Good"

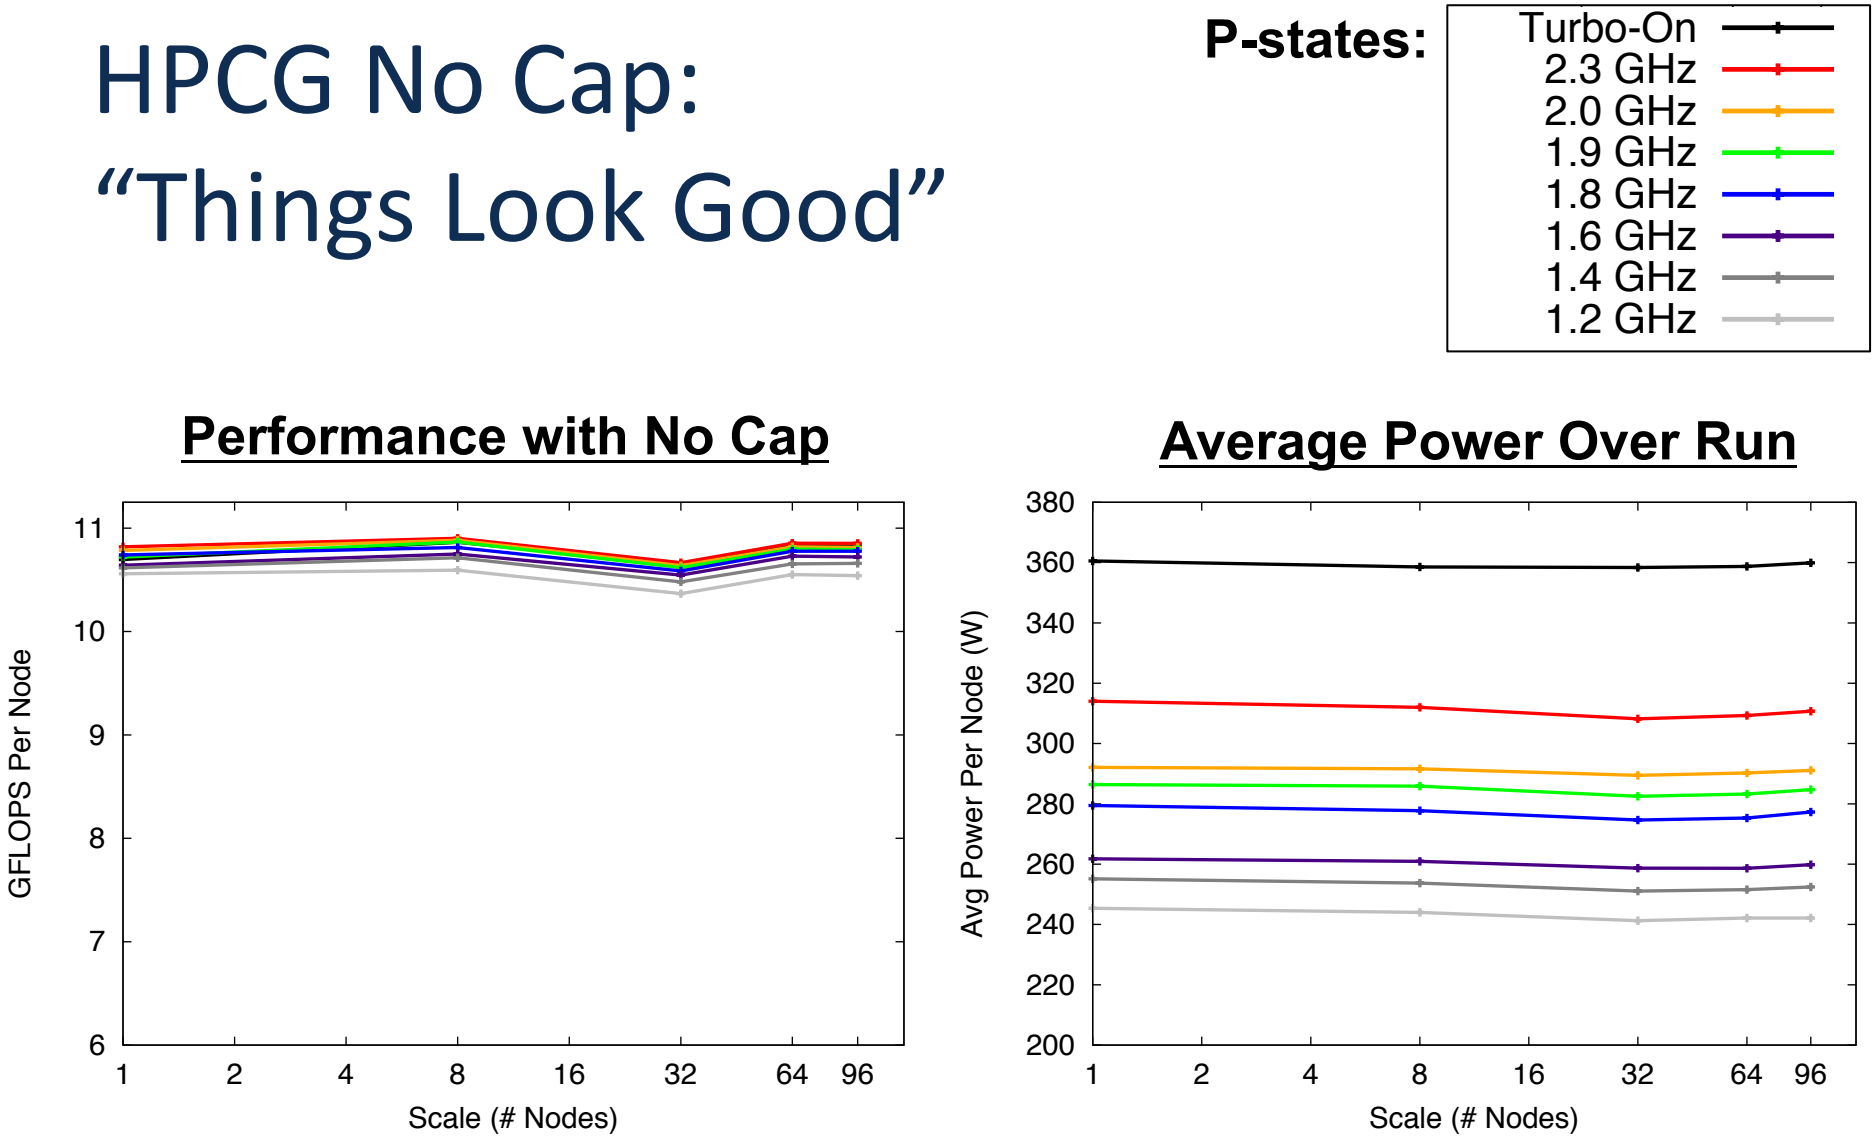

HPCG not affected much by p-state selection;
1.2 GHz: 3% reduced performance, 33% reduced power

HPCG 322W Cap (50%): "Turbo Hits Cap"

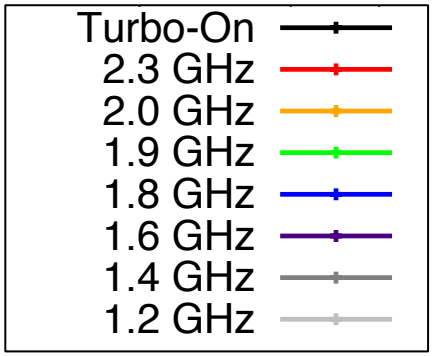

Performance with 322W Cap

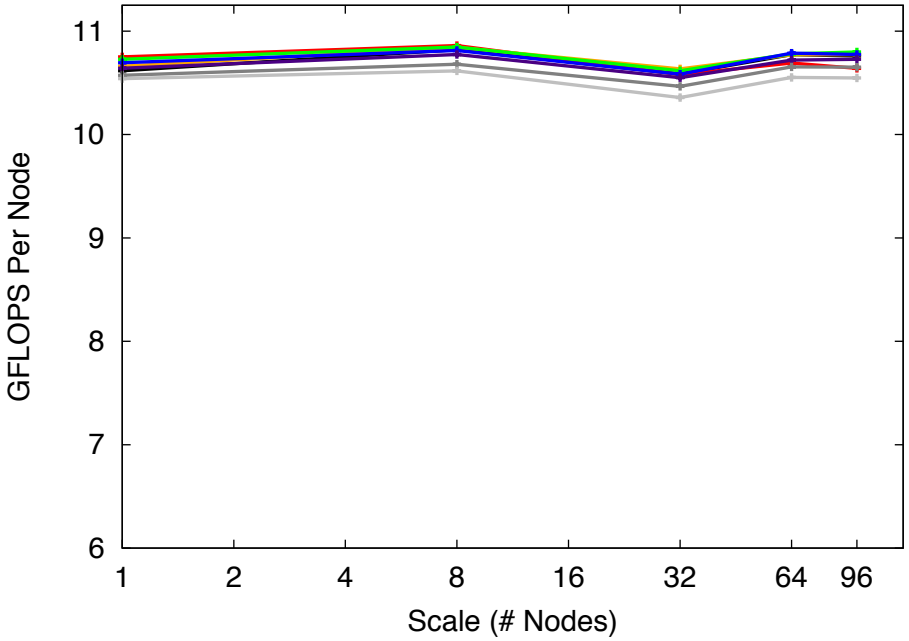

Average Power Over Run

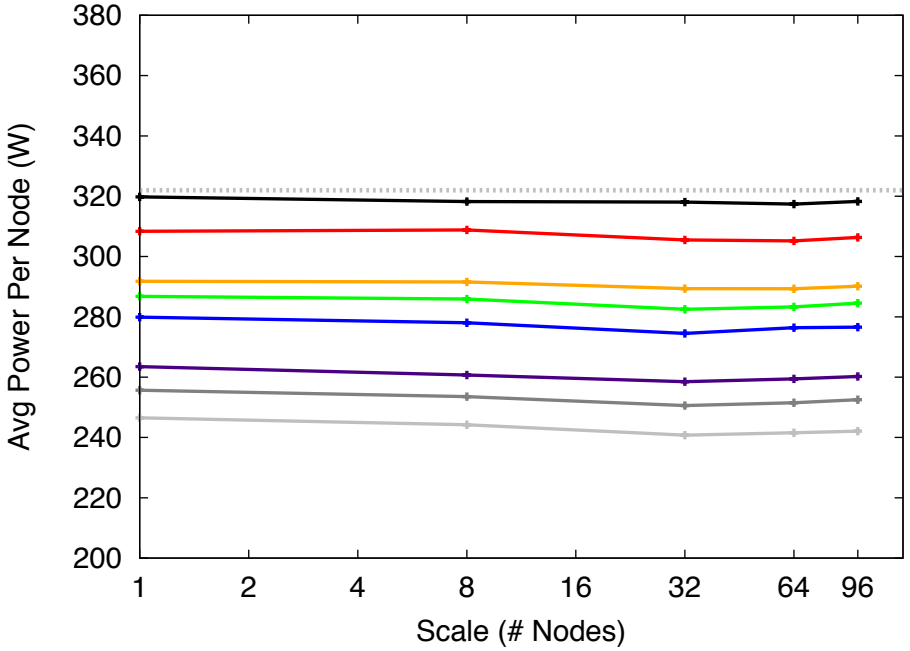

HPCG 230W Cap (0%): "All P-states Hit Cap"

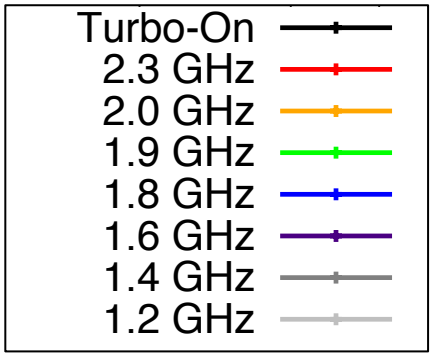

Performance with 230W Cap

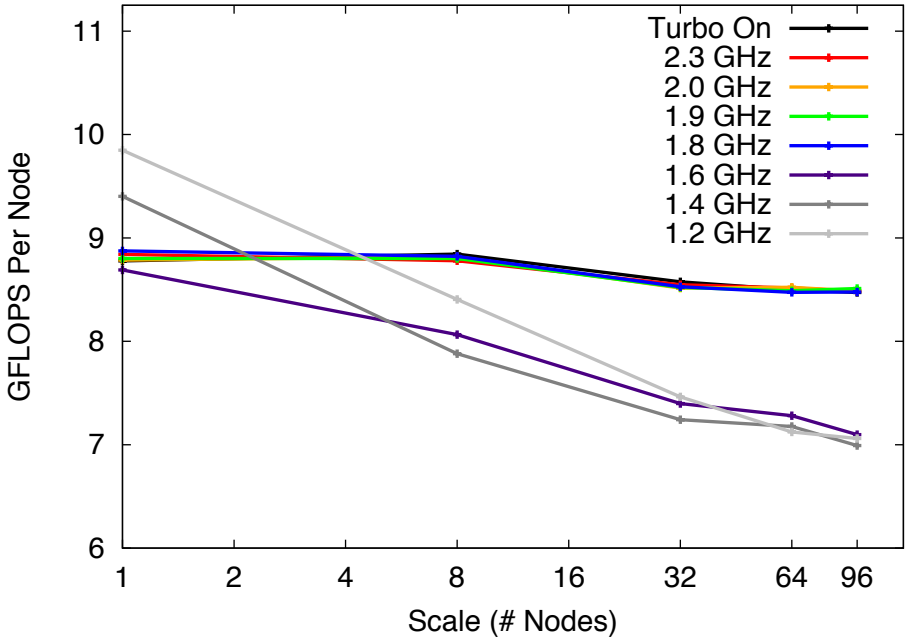

Average Power Over Run

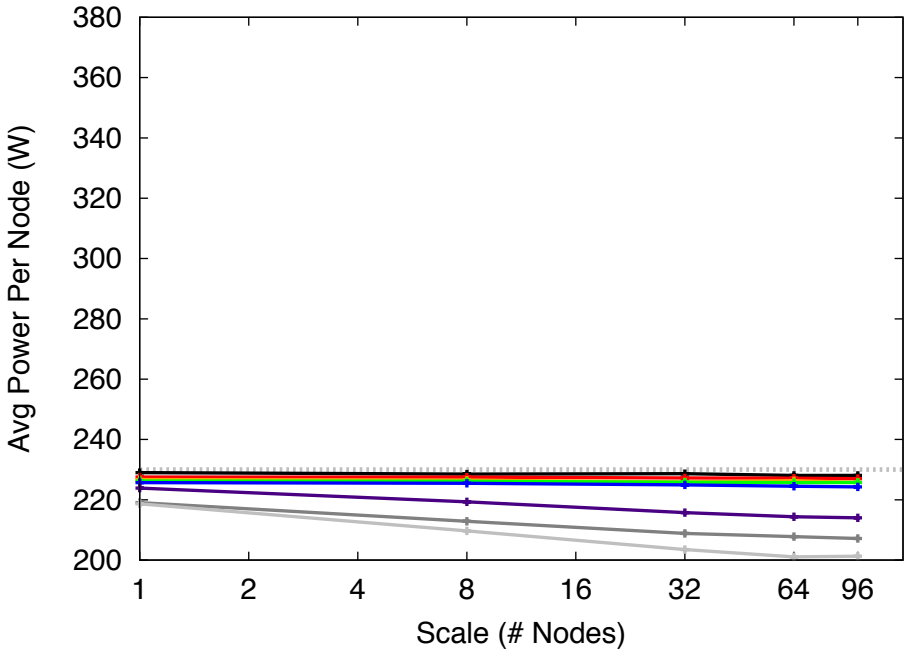

What?
Lowest P-states degrade with scale, not highest!

Summary

Sandia National Laboratories

- S3D makes sense, highest p-states degrade with scale
- CTH looks good, even when hitting cap
- HPL sees degradation with scale, not highest p-states
- HPCG performance degrades with scale for lowest p-states

- Don't know exact reason, still investigating
- Believe to be due to OS-noise like affects, throttling mechanism introducing performance variability across nodes
  - Don't know why performance would drop below 1.2 GHz floor
  - Don't know why not always highest p-states. Maybe some sort of resonance between the application and the power capping policy

Don't Hit the Cap

S3D CPU 0 Frequency Over Time for 10 Nodes Hitting Cap Leads to Frequency Oscillation

Sandia National Laboratories

1.8 Ghz Execution, 322 W Cap (10 of 96 nodes),
Does NOT Hit Cap, All Nodes Run at Same Speed

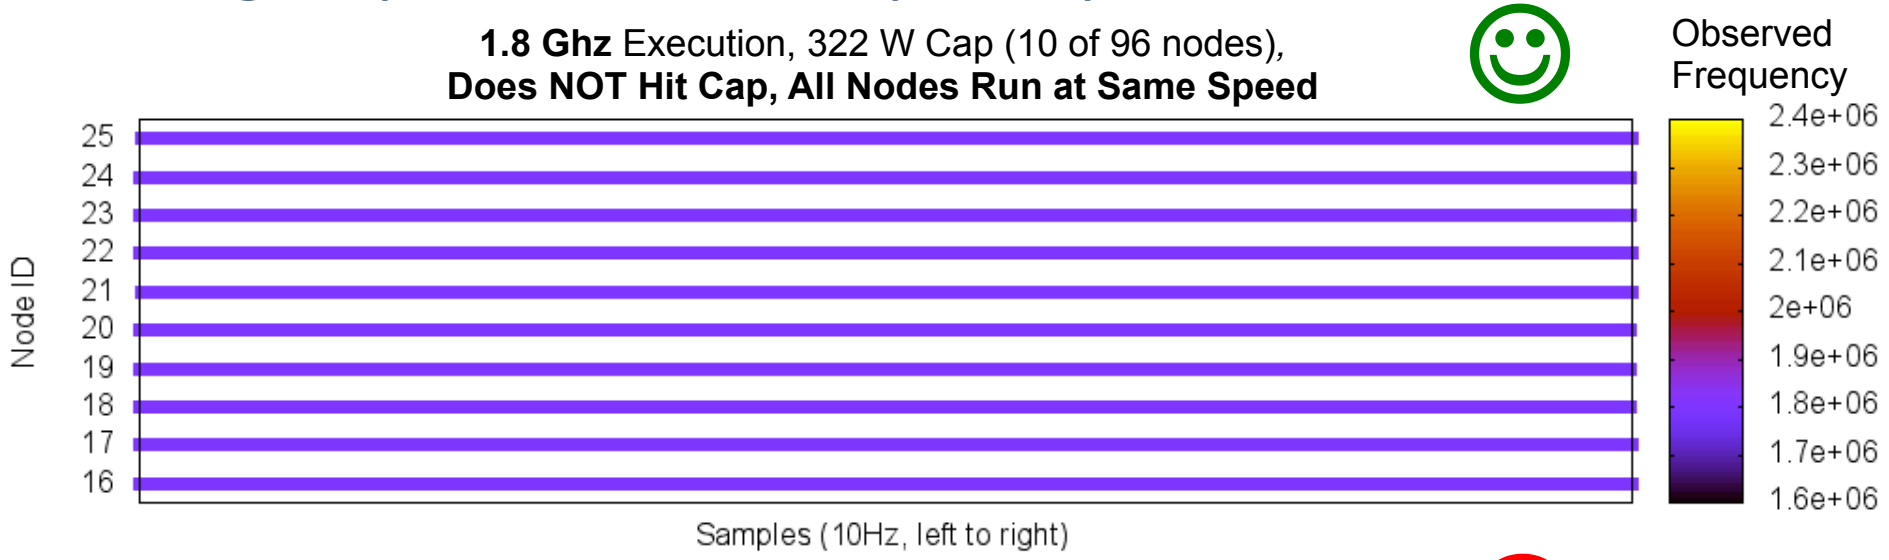

Turbo (Max Speed) Execution, 322 W Cap (10 of 96 nodes),
Cap is Hit Frequently => Performance Variability across Nodes

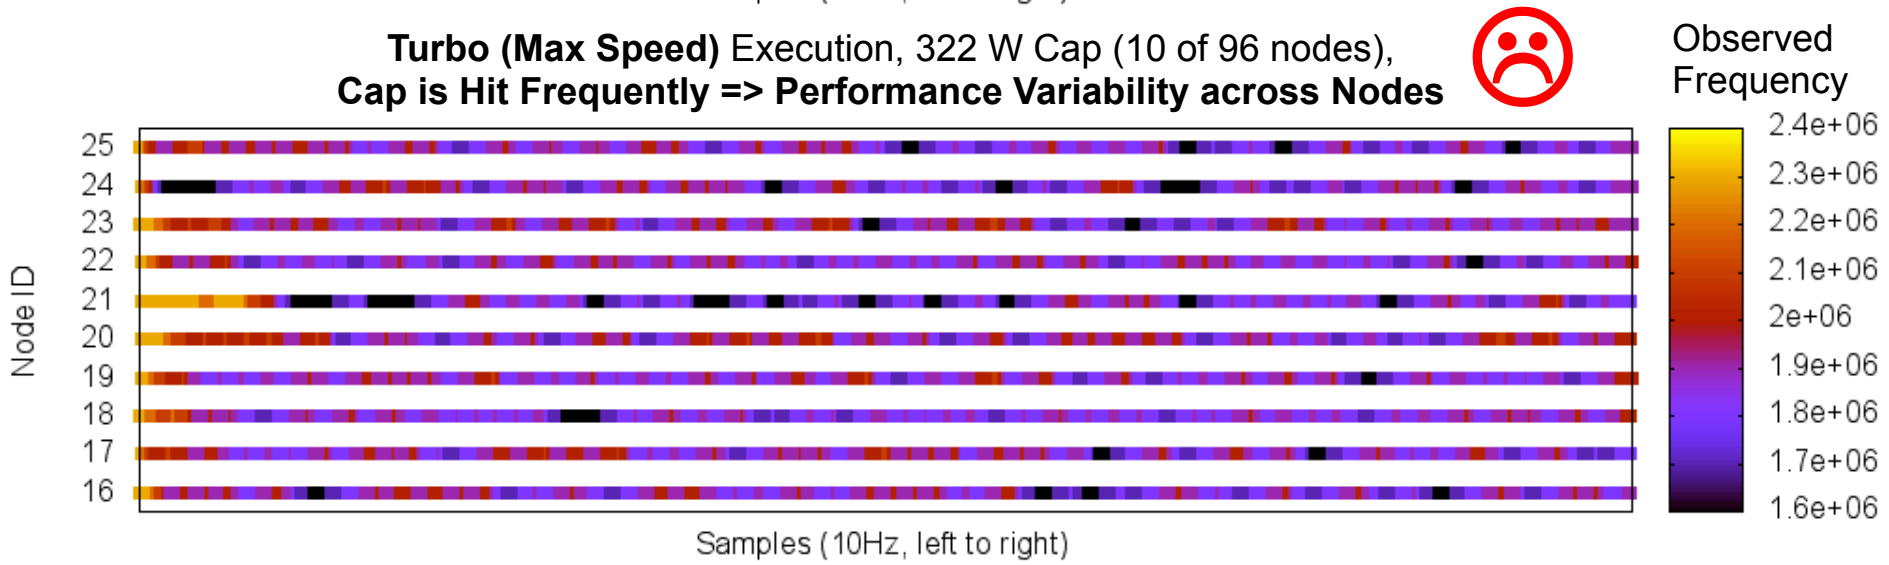

Outline

Sandia National Laboratories

- Background and Motivation
- Cray XC40 Power Management
- Experimental Approach
- Results
  - Don’t hit the power cap
  - Power usage can exceed the cap
- Next Steps and Conclusion

Power Usage Can Exceed the Cap

Sandia National Laboratories

CDFs of 10 Hz node-level power samples taken over entire run,
Power API tool used to collect samples in-band

S3D

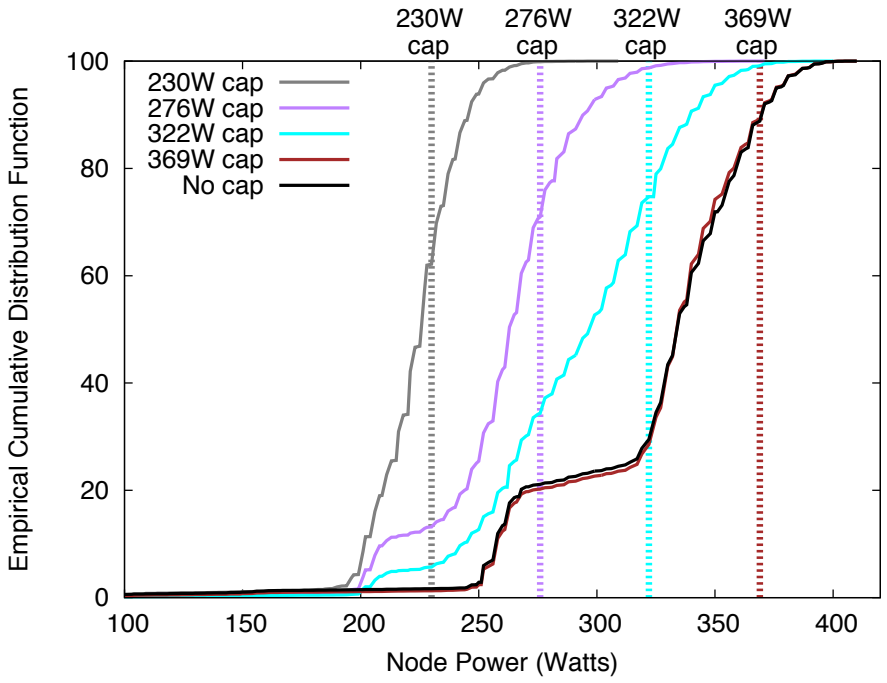

CTH

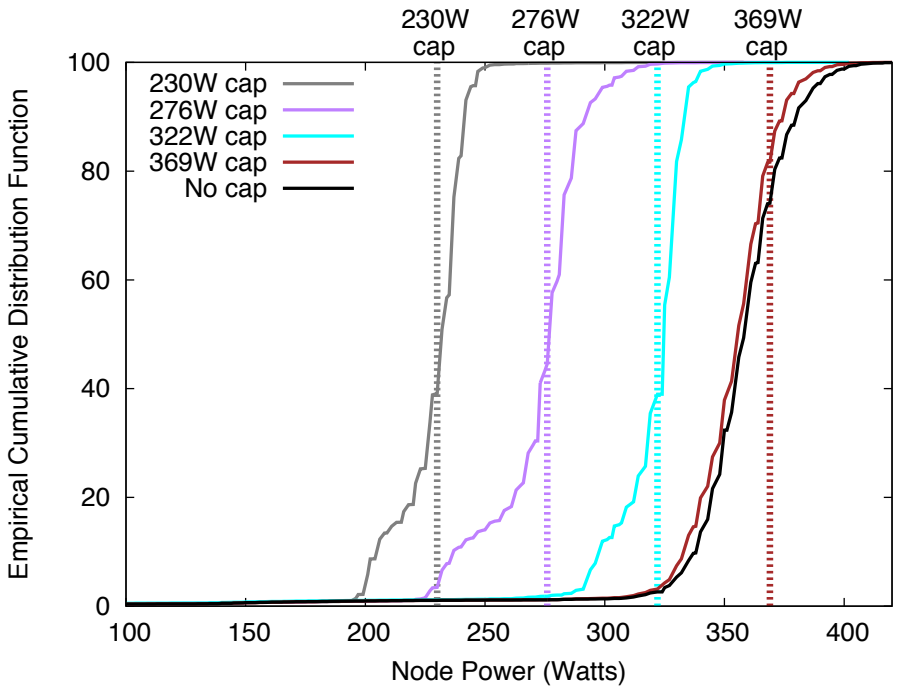

> 50% of samples can exceed cap.
What happens if all nodes exceed cap at same time?

Outline

Sandia National Laboratories

- Background and Motivation
- Cray XC40 Power Management
- Experimental Approach
- Results
  - Don’t hit the power cap
  - Power usage can exceed the cap
- Next Steps and Conclusion

Conclusion

Sandia National Laboratories

- Next steps
  - Further analysis, what exactly is causing poor scalability?
  - Investigate mitigation techniques
  - Dig deeper into what Node Manager is doing and relation to RAPL
  - Test at large-scale on Trinity
- Takeaways, advice to users and system operators
  - Don't hit the node-level cap if possible. Use other power management techniques (e.g., runtime system concurrency throttling) to avoid triggering the capping mechanism's performance throttling.
  - Power usage can exceed the cap for short durations
  - App performance under p-state control better behaved than under power capping. HPCG behaves more similarly to S3D and CTH apps than HPL. (See P-state vs. P-cap comparison in paper)

Acknowledgements

- Ann Gentile
- Jason Repik
- Jim Brandt
- David DeBonis
- Ryan Grant
- James Laros
- Michael Levenhagen
- Steve Martin
- David Rush
- Courtenay Vaughan

PowerAPI BOF
Wed @ 5:30p / Room 13a

Sandia National Laboratories

Thank You

Power Monitoring and Accounting

Sandia National Laboratories

- Power Management Database (PMDB) runs on SMW
- Blade controllers collect power info, funnel to PMDB
  - Maintain energy counter for each node, based on 10 Hz power samples
  - 1 Hz point-in-time power samples for each node stored in PMDB
- Users have in-band access to Cray node-level power and energy counters: **`/sys/cray/pm_counters/energy`**
- Cray Resource Utilization Reporting (RUR) gathers energy used for each job run. Report written to user's ~/rur.out file:

```
uid: 20983, apid: 175244, jobid: 349.mutrino-moab,
cmdname: ./bin/mpicth_244, plugin: energy
{"nodes_throttled": 2, "nodes_with_changed_power_cap": 0,
"max_power_cap_count": 0, "energy_used": 25015991,
"max_power_cap": 0, "nodes_accel_power_capped": 0,
"min_power_cap": 0, "min_power_cap_count": 0,
"nodes_power_capped": 0, "nodes": 96}
```

Power Caps – Performance Impact

Sandia National Laboratories

- As power budget is decreased, application power hits the cap
  - System slows down the processor to stay under the cap
  - Performance degrades
- By running at a lower frequency, avoid hitting the cap
  - With 96 nodes & 230W cap, better at 1.4 Ghz than at 2.3Ghz
  - Best is the highest frequency that stays below the cap
    - Depending on node count and power cap level
    - With 96 nodes & 230W cap, best choice is 1.8 GHz
    - With 96 nodes & 276W cap, best choice is 2.0 GHz
- Why the degradation?
  - System slows down processors by varying amounts, at different times
  - Induces load imbalance

Relative Performance

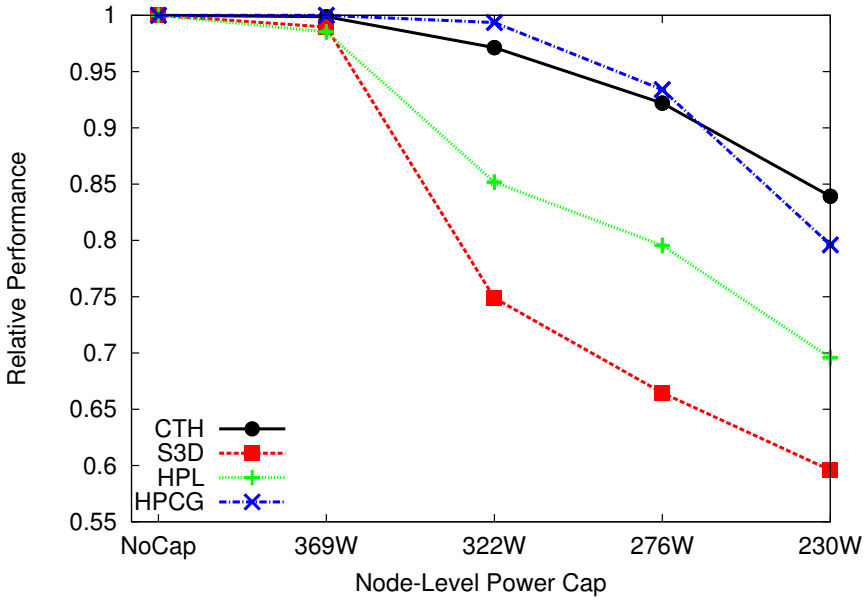

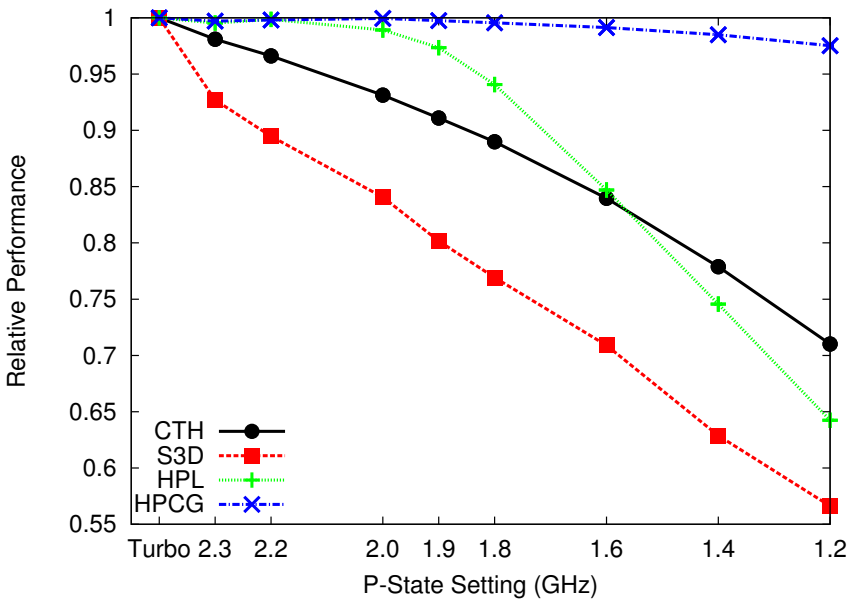

Relative Power

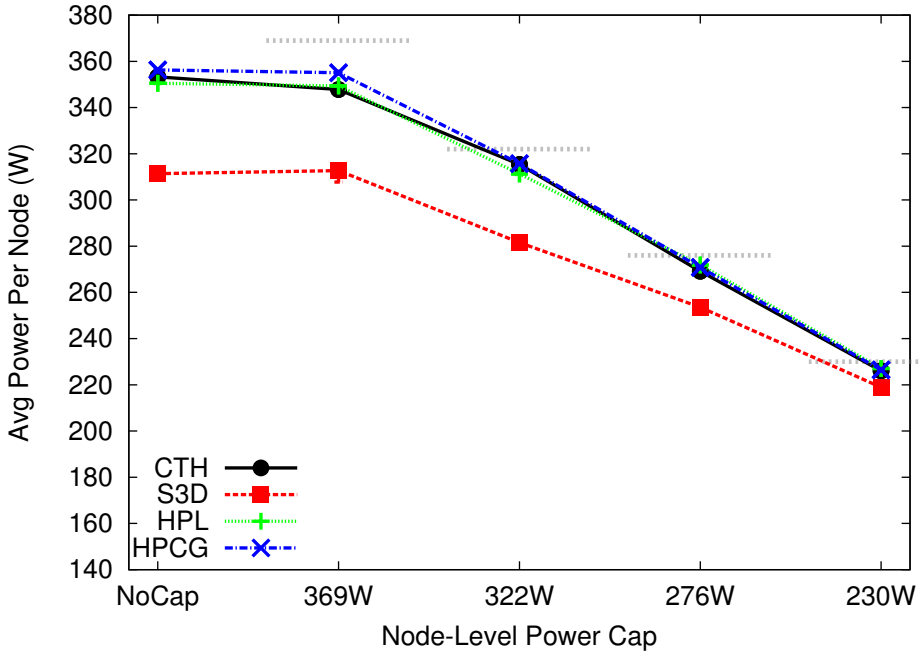

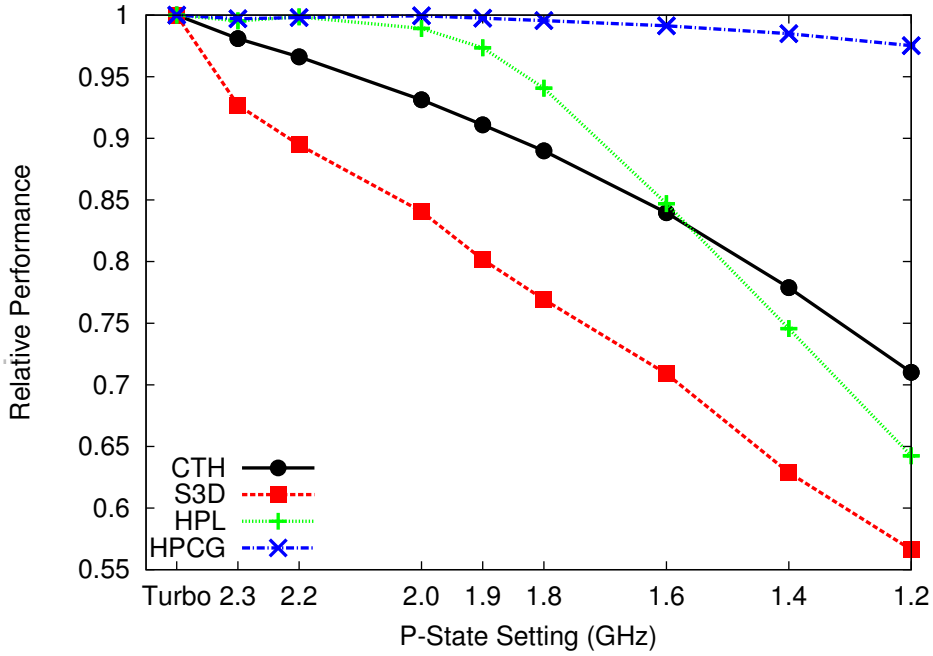

Relative Energy

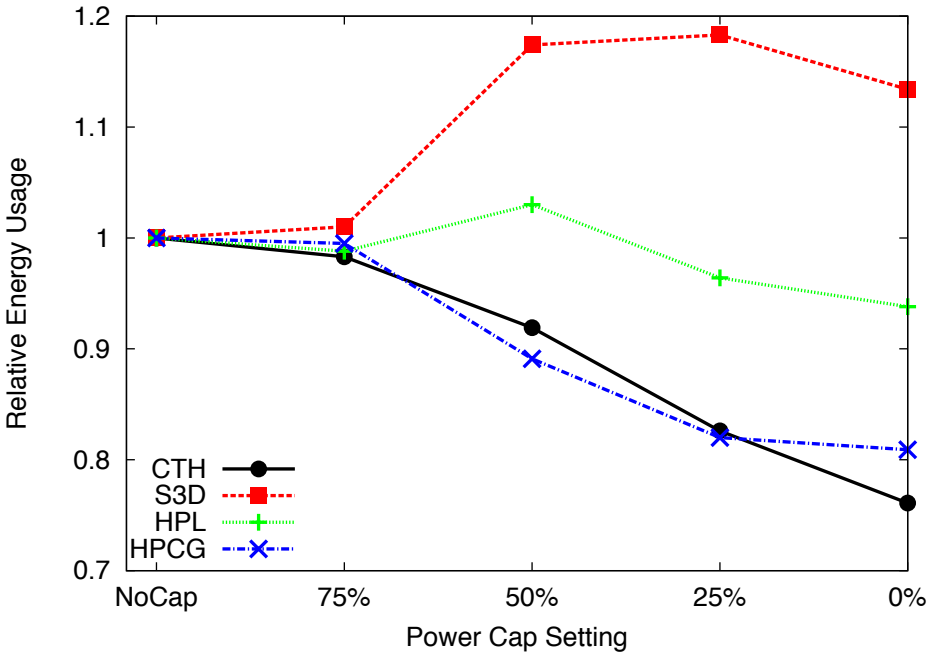

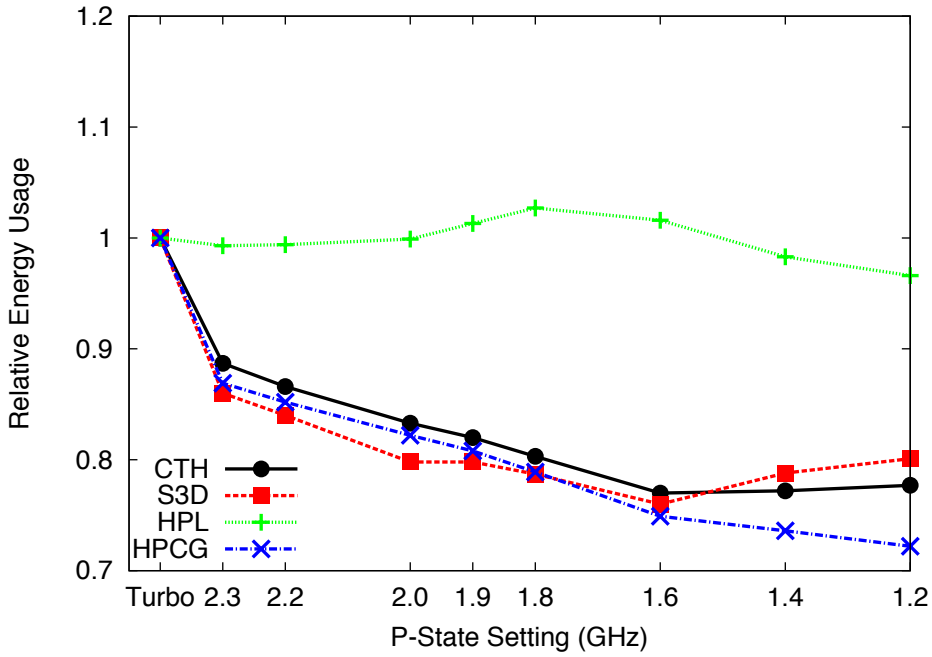

Relative Energy Efficiency