MANUFACTURING VARIABILITY IN HPC

Tenth Annual Workshop for the Energy Efficient HPC Working Group

SC19 - Denver

PANEL AGENDA

- **Manufacturing Variation and Other Challenges to Platform Level Monitoring and Management**

- **Speakers:**
  - Rabin Sugumar, Marvell
  - Barry Rountree, LLNL
  - Thomas Ilsche, TU-Dresden
  - Steven Martin, Cray-HPE (Moderator)

- Group Q&A

QUESTIONS TO FRAME THIS SESSION

- What is manufacturing variability, and why does it happen
- Manufacturing variability experiences, or challenges observed
- Assuming this issue is not going away
  - How does an HPC site adapt to it?
  - Do you leverage software to manage it?

MARVELL®

Manufacturing Variability

Rabin Sugumar

Chief Architect

Server Processors

Nov 2019

Overview of Talk

- How are processors shipped
- Sources of variation in manufacturing
- How are processors different part to part wrt power
- Other sources of variation
- Turbo algorithms
- What can be done to mitigate

© 2018 Marvell Confidential, All Rights Reserved.

How are processors shipped

- Each part goes through testing (ATE and SLT) for functionality
- Parts that pass functional test undergo power/frequency tests
  - Does it run at base and max turbo frequency for SKU ?
  - Does it run within TDP at base frequency for SKU ?
- Depending on volumes required at point in time binning may be more or less precise. Example:
  - When most volume is for TDP=180W, all parts under 180W shipped
  - When equal quantities needed at TDP=150W and TDP=180W more precise binning to separate parts below 150W from parts below 180W
- Less precise binning allows more turbo headroom
  - Hardware vendor is providing better parts on average
  - But results in more variation

© 2018 Marvell Confidential, All Rights Reserved.

Sources of Variation in Manufacturing

- Patterning proximity effects
- Line edge roughness, line width roughness
- Polish variations
- Gate dielectric variations
- Random dopant fluctuation
- Variation in anneals
- Variation in implants
- And others …

- Continued effort in manufacturing to reduce variation
- Variations unlikely to go away or even reduce significantly

Ref: K.J.Kuhn et al., “Process Technology Variation”, 2011

© 2018 Marvell Confidential, All Rights Reserved.

Power Variation among Parts

- Leakage current – close to 3x variation among parts
- Vmin for a target frequency – greater than 50mV variation
  - Lower the Vmin, lower the power – quadratic effect
- Dynamic power dissipated
- These metrics are correlated – for instance leakier parts may operate at lower Vmin
- Across large die behavior of cores will be different
  - Different Vmin, different leakage/dynamic power

© 2018 Marvell Confidential, All Rights Reserved.

Leakage Variation

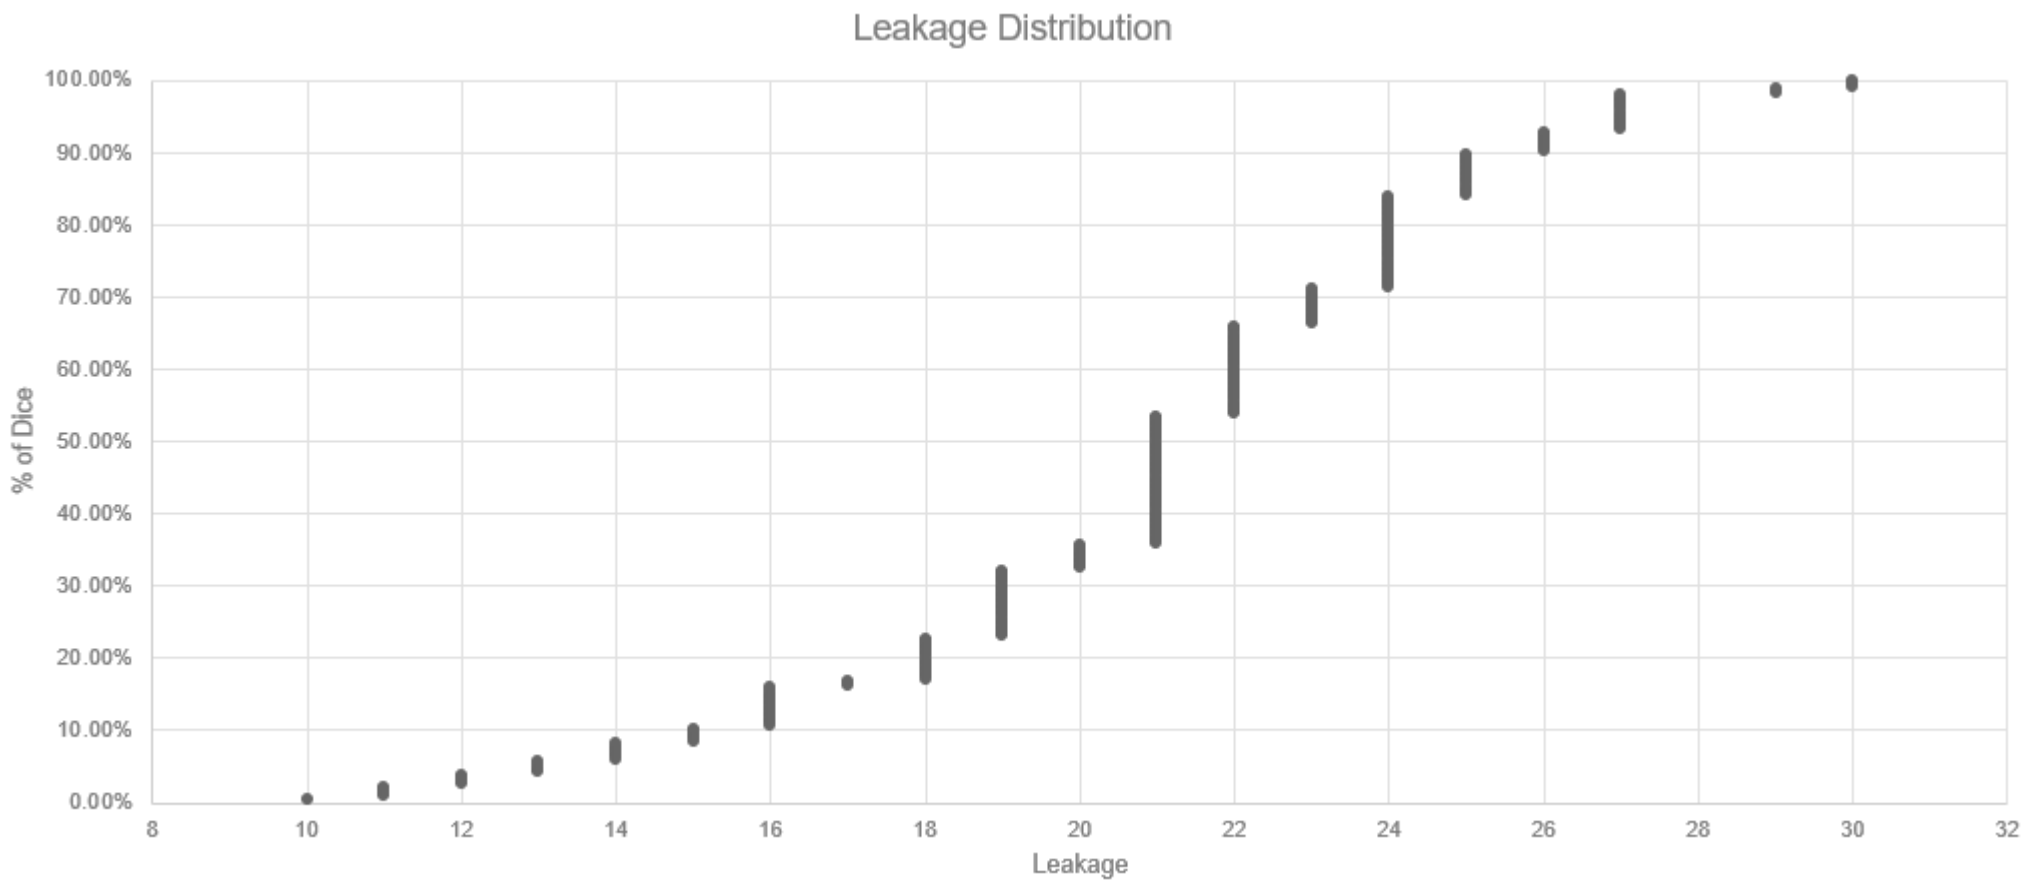

© 2018 Marvell Confidential, All Rights Reserved.

Vmin Variation

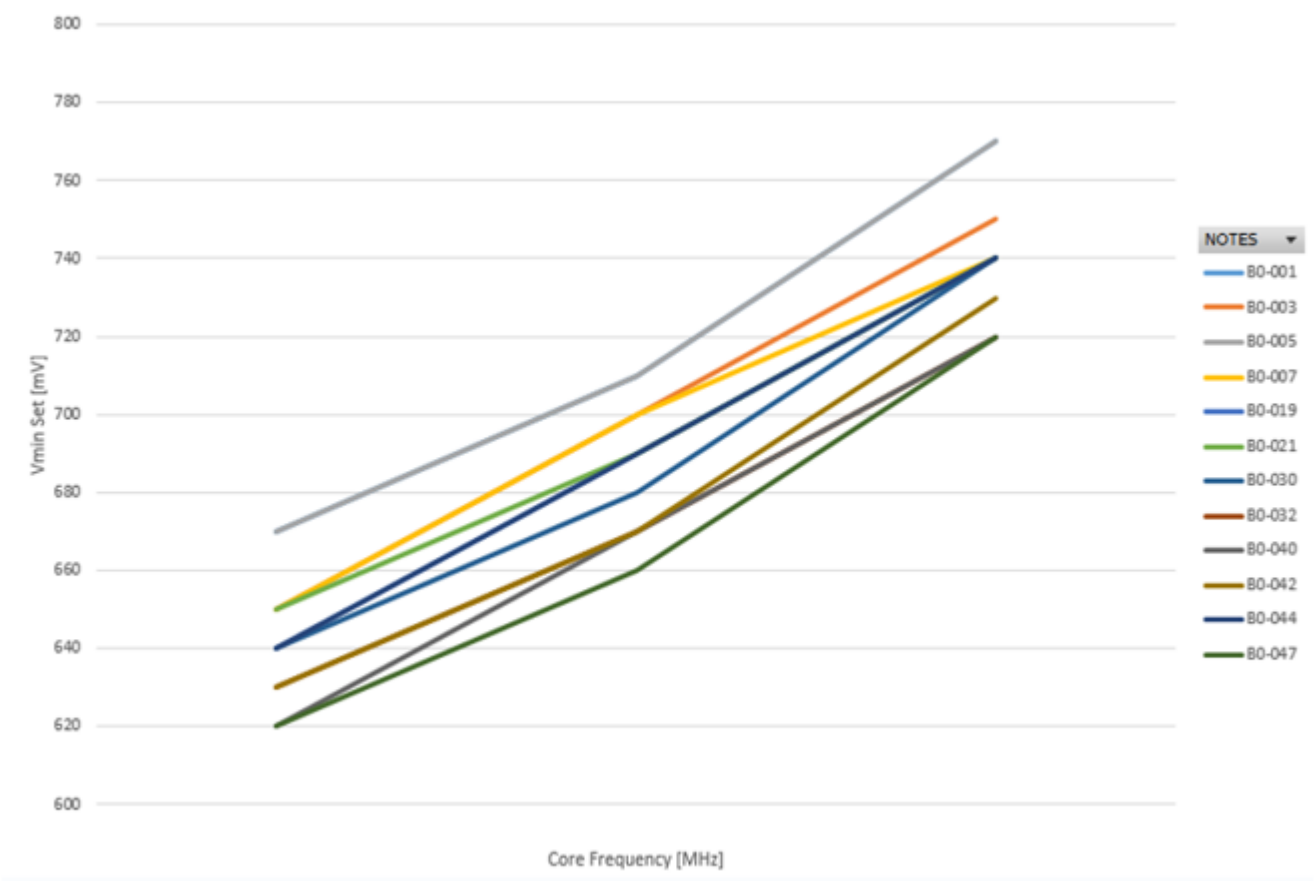

© 2018 Marvell Confidential, All Rights Reserved.

Chip Power Variation

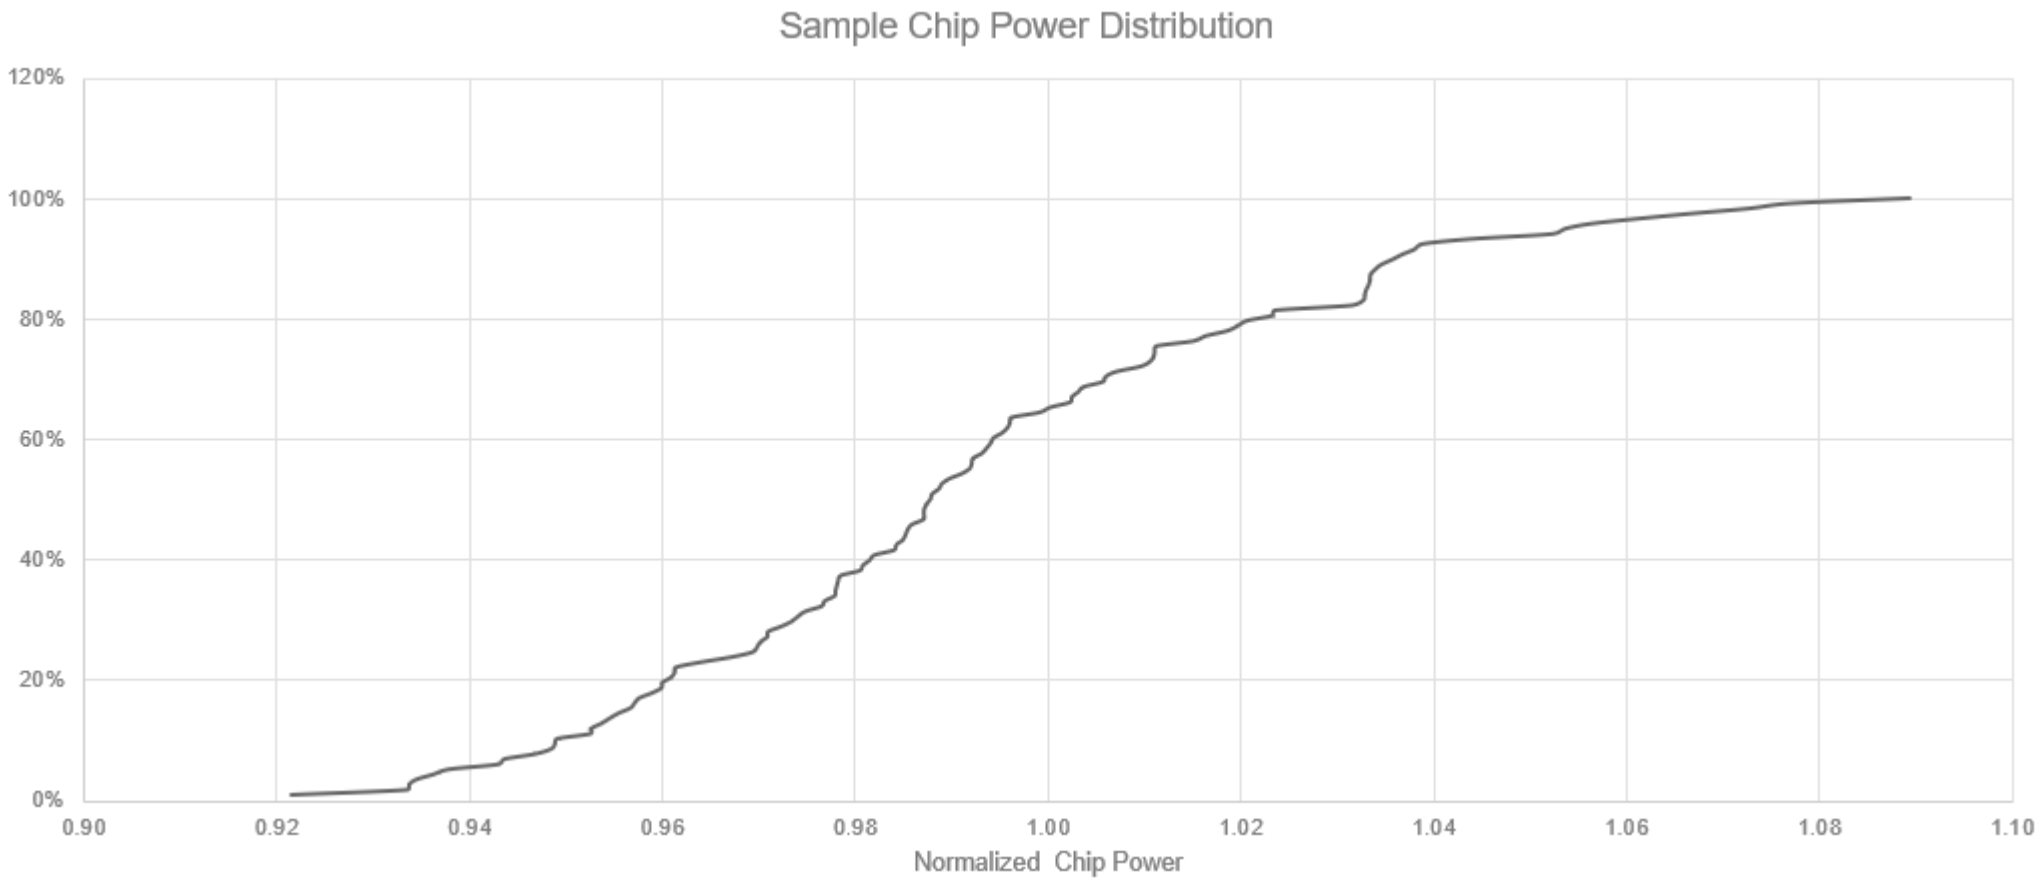

© 2018 Marvell Confidential, All Rights Reserved.

Other sources of variation

- Temperature
  - Temperature across large systems is not uniform
  - Power dissipation higher at higher temperature (higher leakage)
- Software/Data
  - Different nodes in a system may be executing different code/data
  - Higher IPC results in more power (greater activity)
  - Higher data variation results in more power (greater bus activity)
- Wear
  - Device characteristics change over time
  - Platform characteristics change over time
  - For various reasons certain nodes may have more wear than others

© 2018 Marvell Confidential, All Rights Reserved.

Turbo Algorithm

- Intelligent on die controller that varies frequency and voltage to maximize performance within power constraints
- At its simplest
  - Monitors power (sample periodically)
  - When power is within TDP raises frequency/voltage of active cores
  - When power rises above TDP lowers frequency/voltage of active cores
- Optimization focus tends to be localized
  - Maximize performance of a die or performance of a core
  - Usually doesn’t consider other die in the larger system

© 2018 Marvell Confidential, All Rights Reserved.

Is variation bad ?

- Is the ask that all die run at performance level of best die or
- Is the ask that all die run at performance level of worst die

© 2018 Marvell Confidential, All Rights Reserved.

Reduce Variation

- Disable turbo – always run at base frequency
  - Problem: Base frequency may be too conservative for many workloads
- Specify a frequency for a workload to bin
  - Problem 1: May work for certain large customers with well defined workloads, but not a general solution
  - Problem 2: Expensive
- Make turbo algorithm system aware. E.g., set max performance level of all cores/die to worst turbo core/die participating in an application run
  - Problem: Requires firmware/hardware networking and algorithms
  - Doable with time and effort

© 2018 Marvell Confidential, All Rights Reserved.

Handle Variation at Workload Level

- Rewrite algorithms to exploit some cores running faster than others
- Create subclusters of nodes that have similar behavior

© 2018 Marvell Confidential, All Rights Reserved.

Conclusions

- Several sources of die performance/power variation
- Variations unlikely to go away
- Will be good to find ways to exploit performance of better die
- Turbo algorithms could become more system aware

© 2018 Marvell Confidential, All Rights Reserved.

Population studies in processor variation: Broadwell and Sandy Bridge processors

CASC
Center for Applied Scientific Computing

15-Nov-2019

Barry Rountree
rountree@llnl.gov

LLNL-PRES-XXXXXX
This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under contract DE-AC52-07NA27344. Lawrence Livermore National Security, LLC

Lawrence Livermore National Laboratory

Broadwell Variation

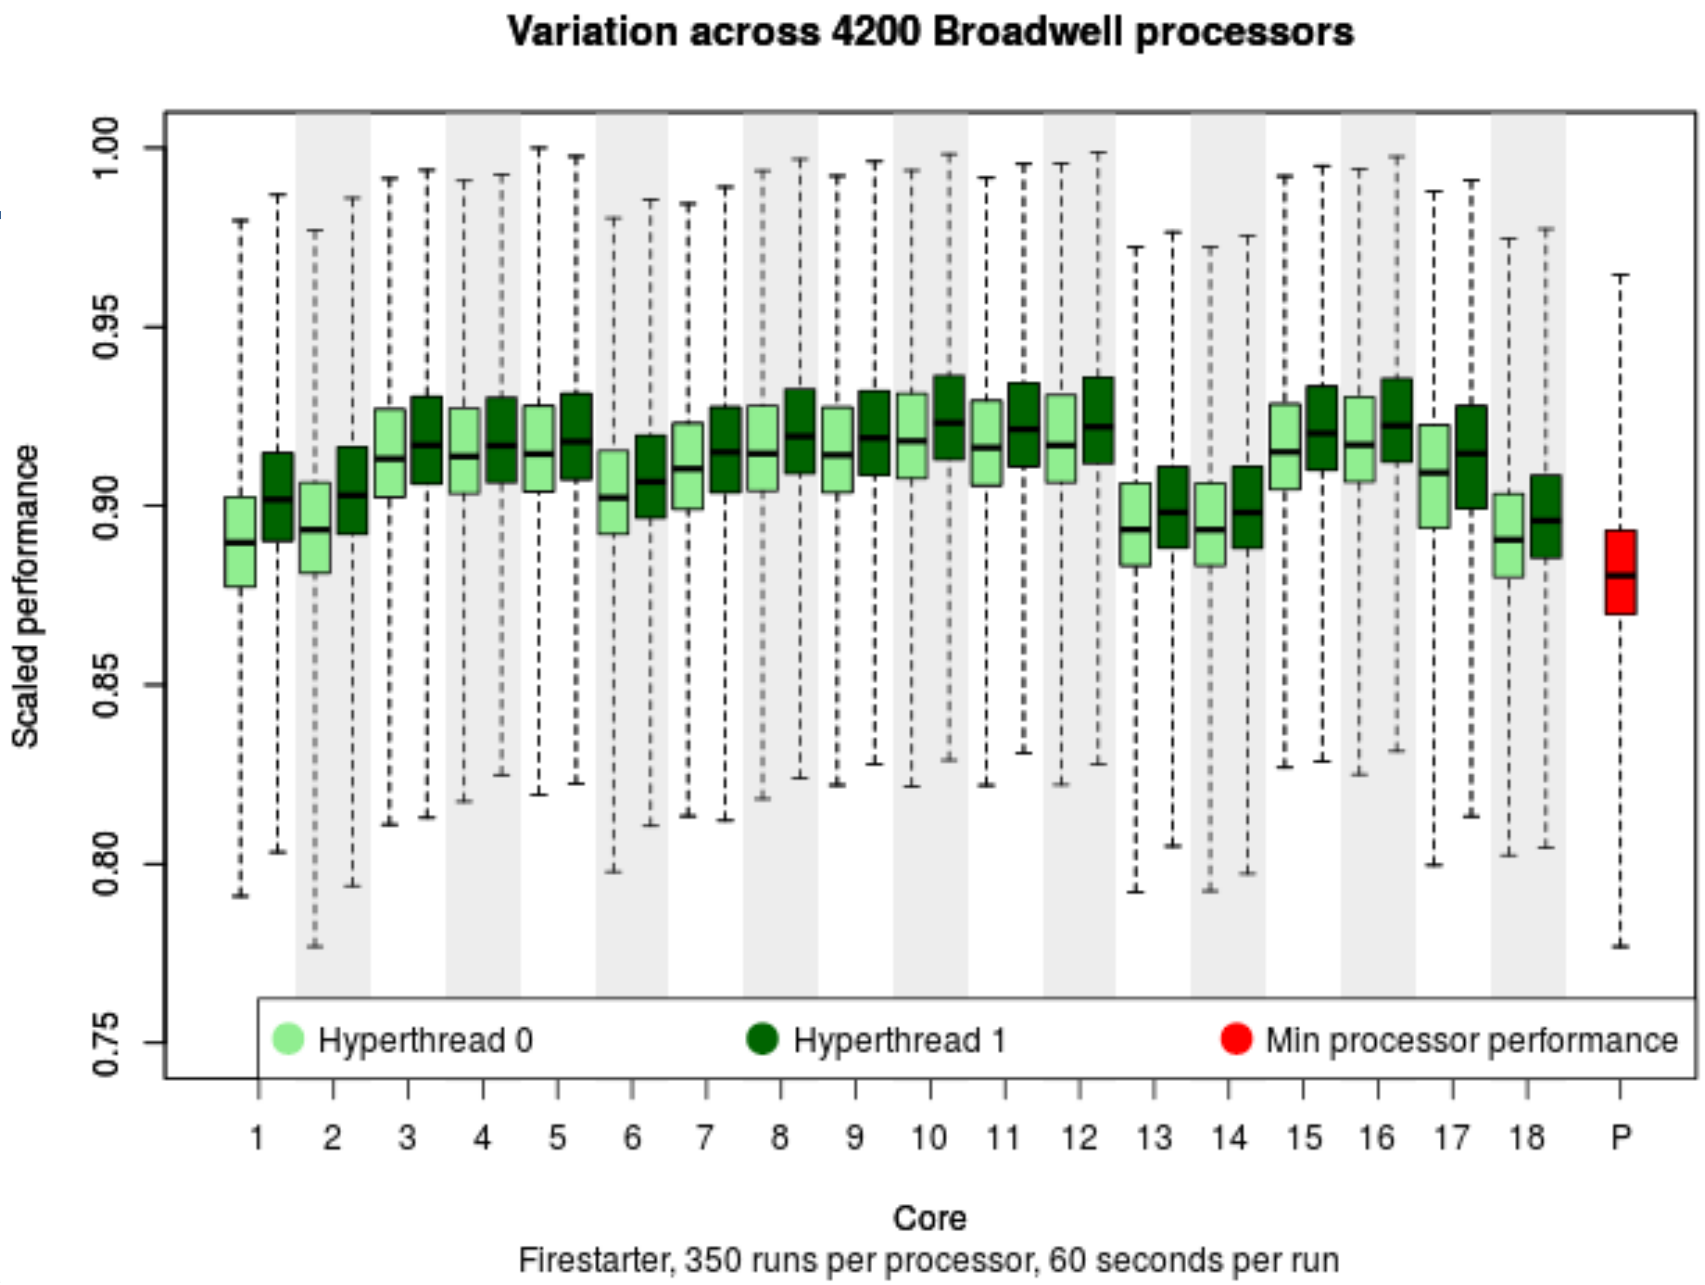

- Jade
  - 4200 Broadwell processors
- Firestarter
  - Embarrassing parallel, AVX-heavy
- Experiment
  - 60 seconds per run
  - 350 runs per processor
  - Performance measured per hyperthread
  - No power bounds
  - Turbo enabled
- Results
  - 50% of values in green boxes
  - 20% variation

Broadwell Variation

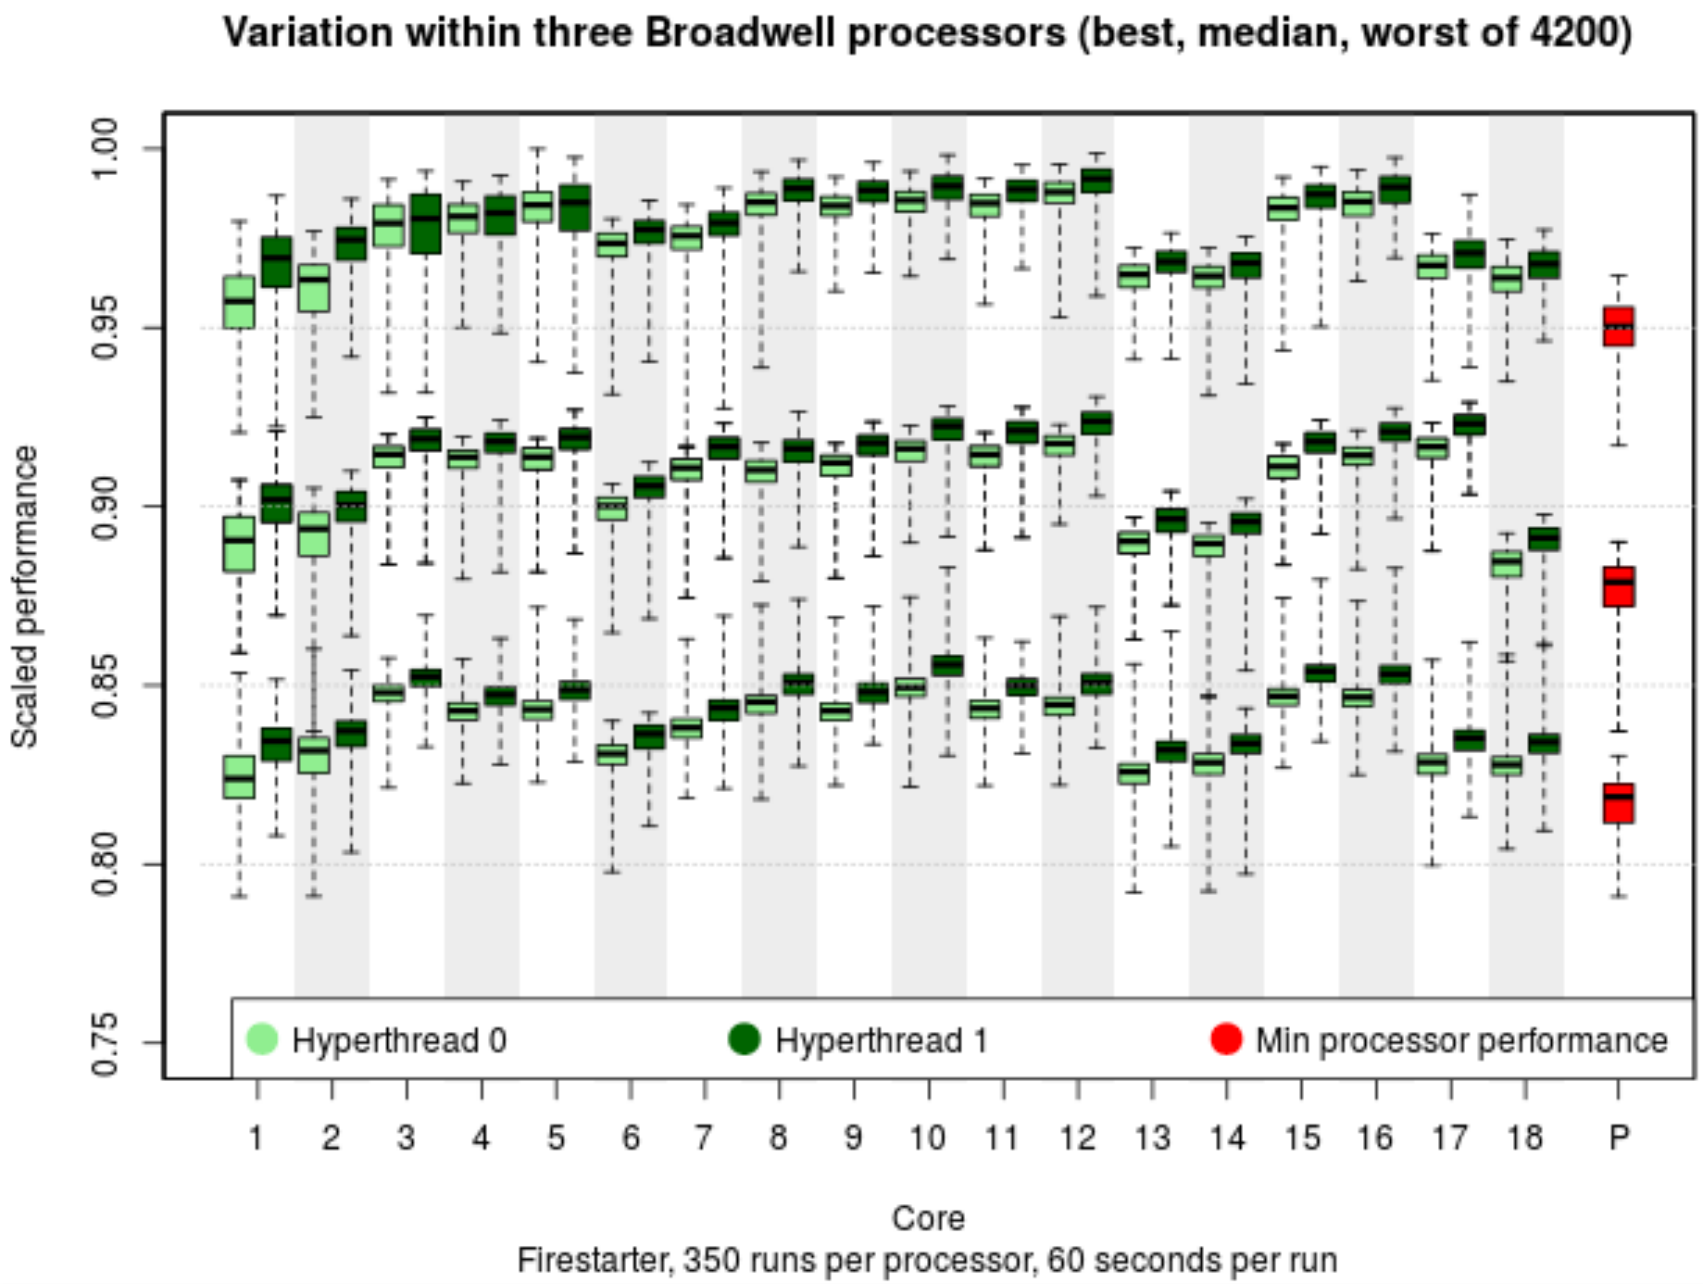

- Jade
  - Best, worst and median processors
- Results
  - 50% of values in green boxes
  - Run-to-run variation
  - Processor-to-processor variation
  - Core-to-core variation
  - Hyperthread-to-hyperthread variation

Sandy Bridge variation across four power bounds

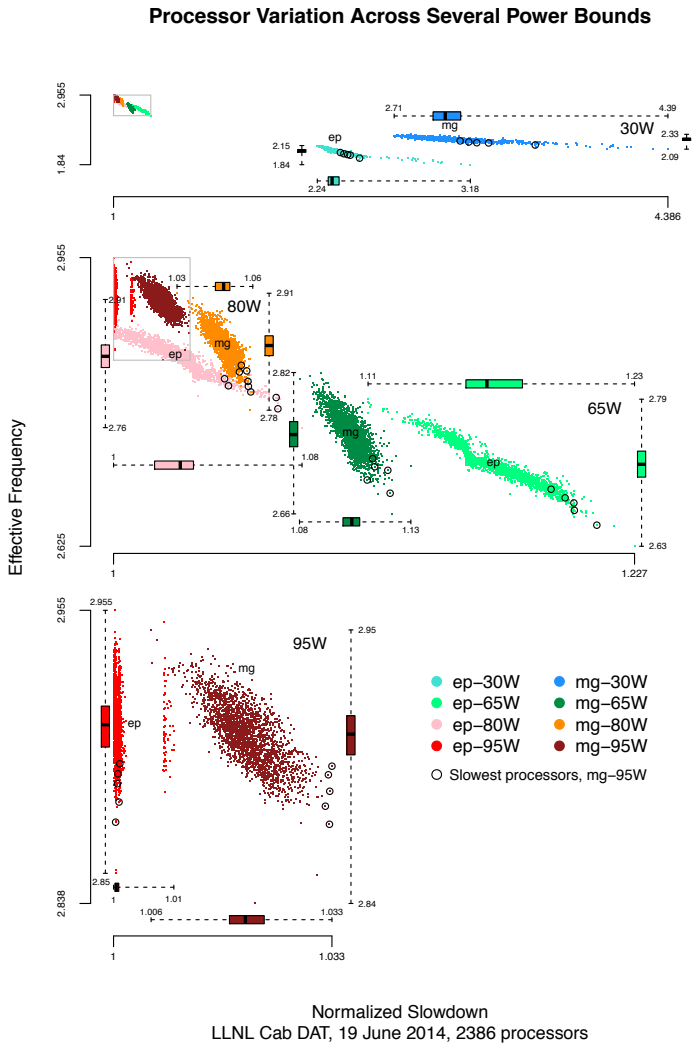

- Cab cluster
  - 2386 Sandy Bridge processors
- Experiments
  - NAS EP, MG
  - Single-node, no communication
  - Slowdown normalized to fastest uncapped run
  - Five worst MG 95W processors circled in each experiment
- Results
  - Performance degradation depends in part on sensitivity to CPU clock speed
  - Power capping exacerbates variation

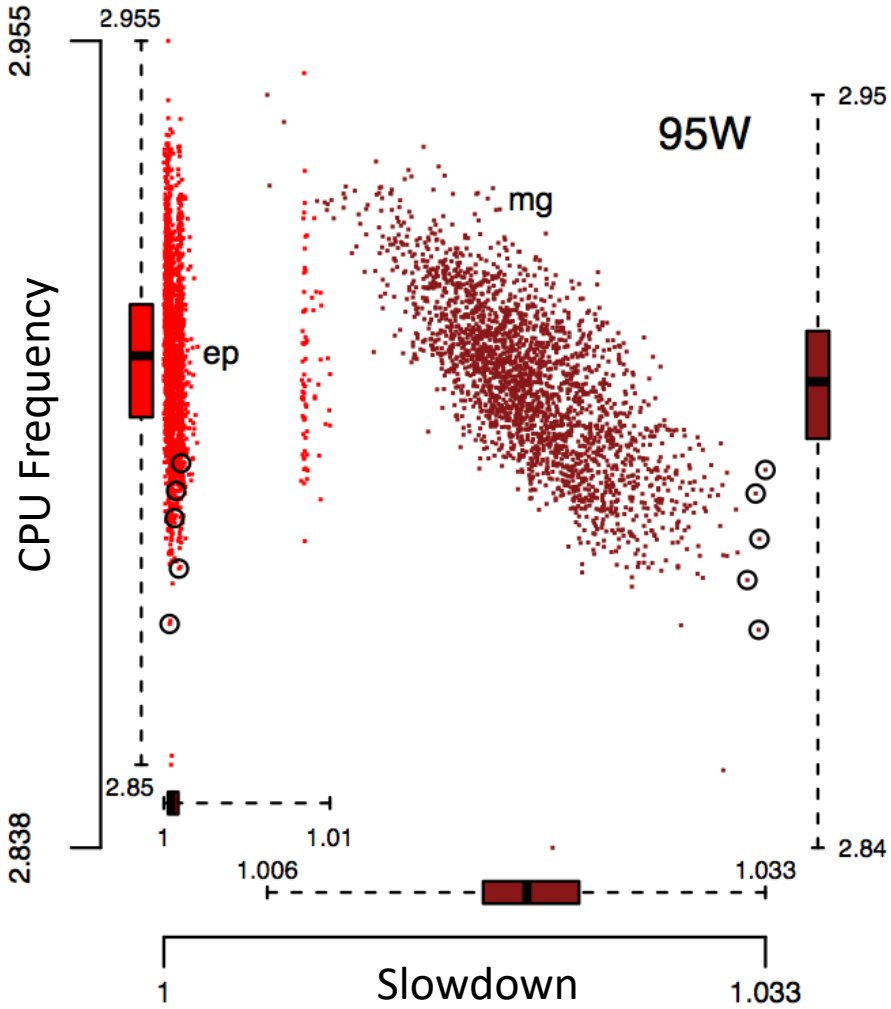

Sandy Bridge variation across four power bounds

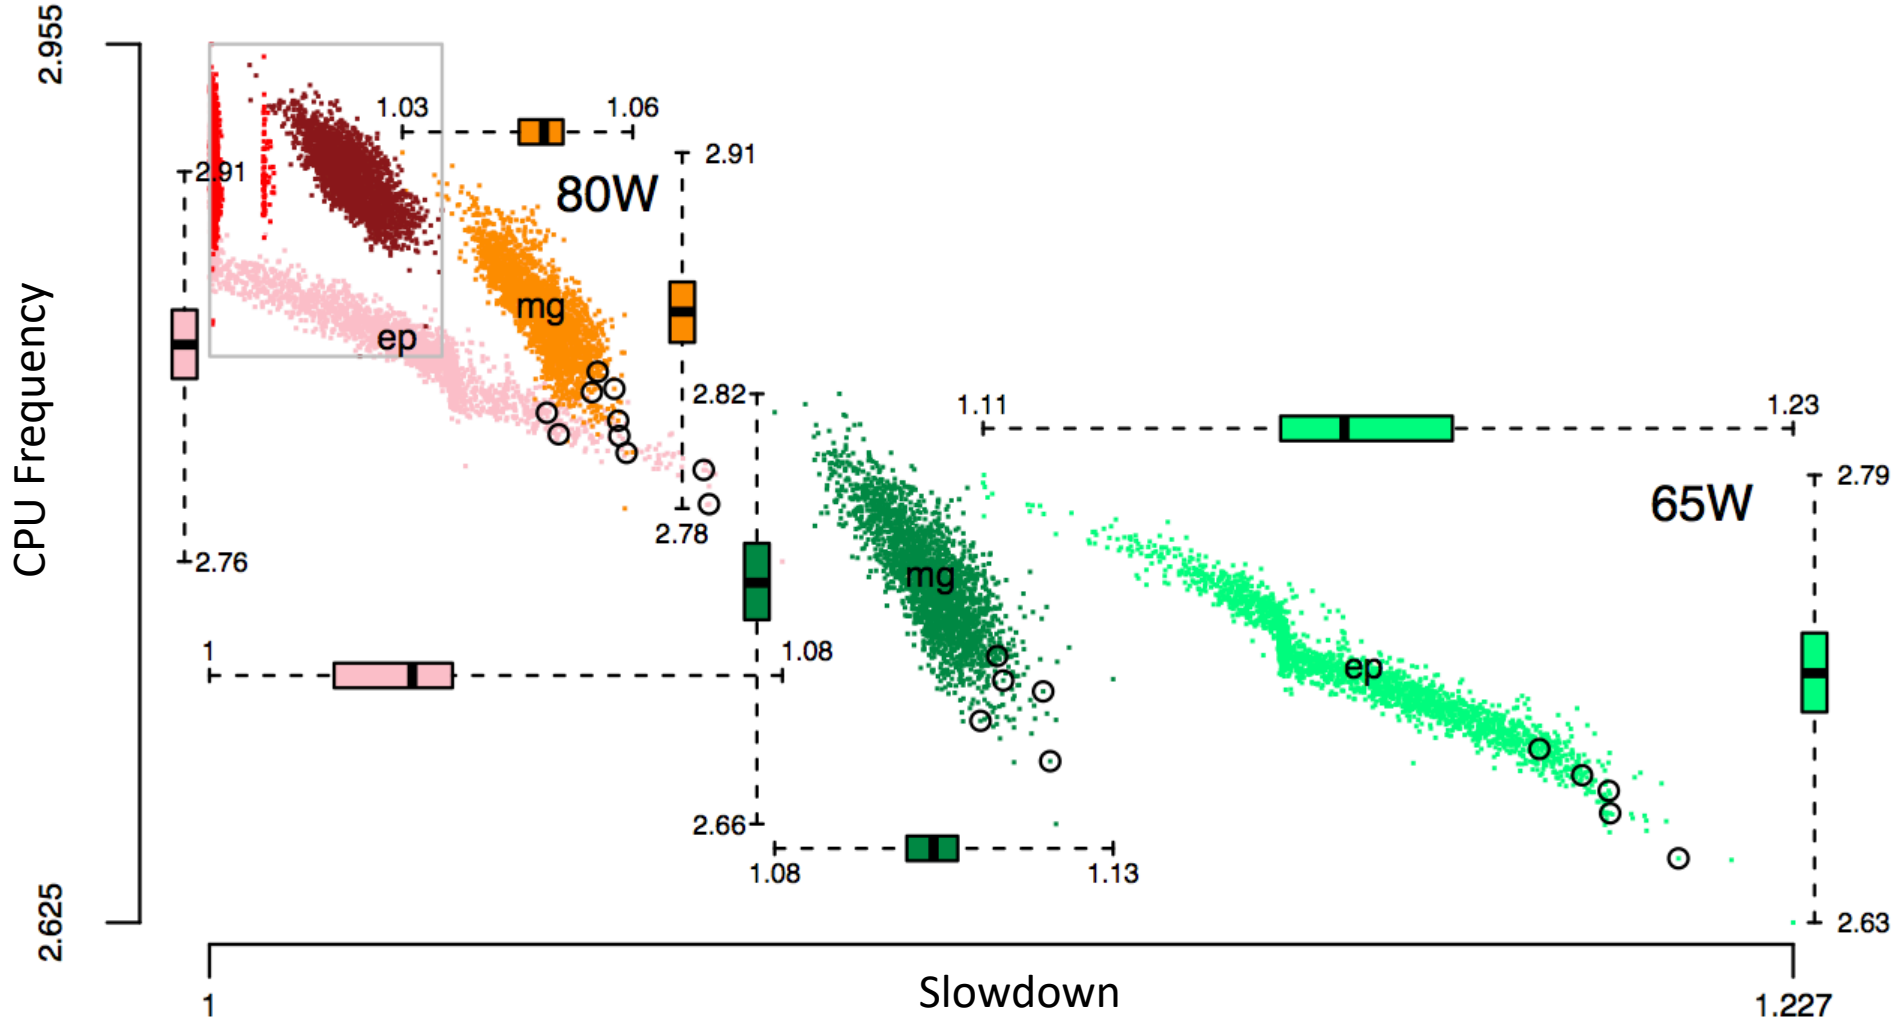

Sandy Bridge variation across four power bounds

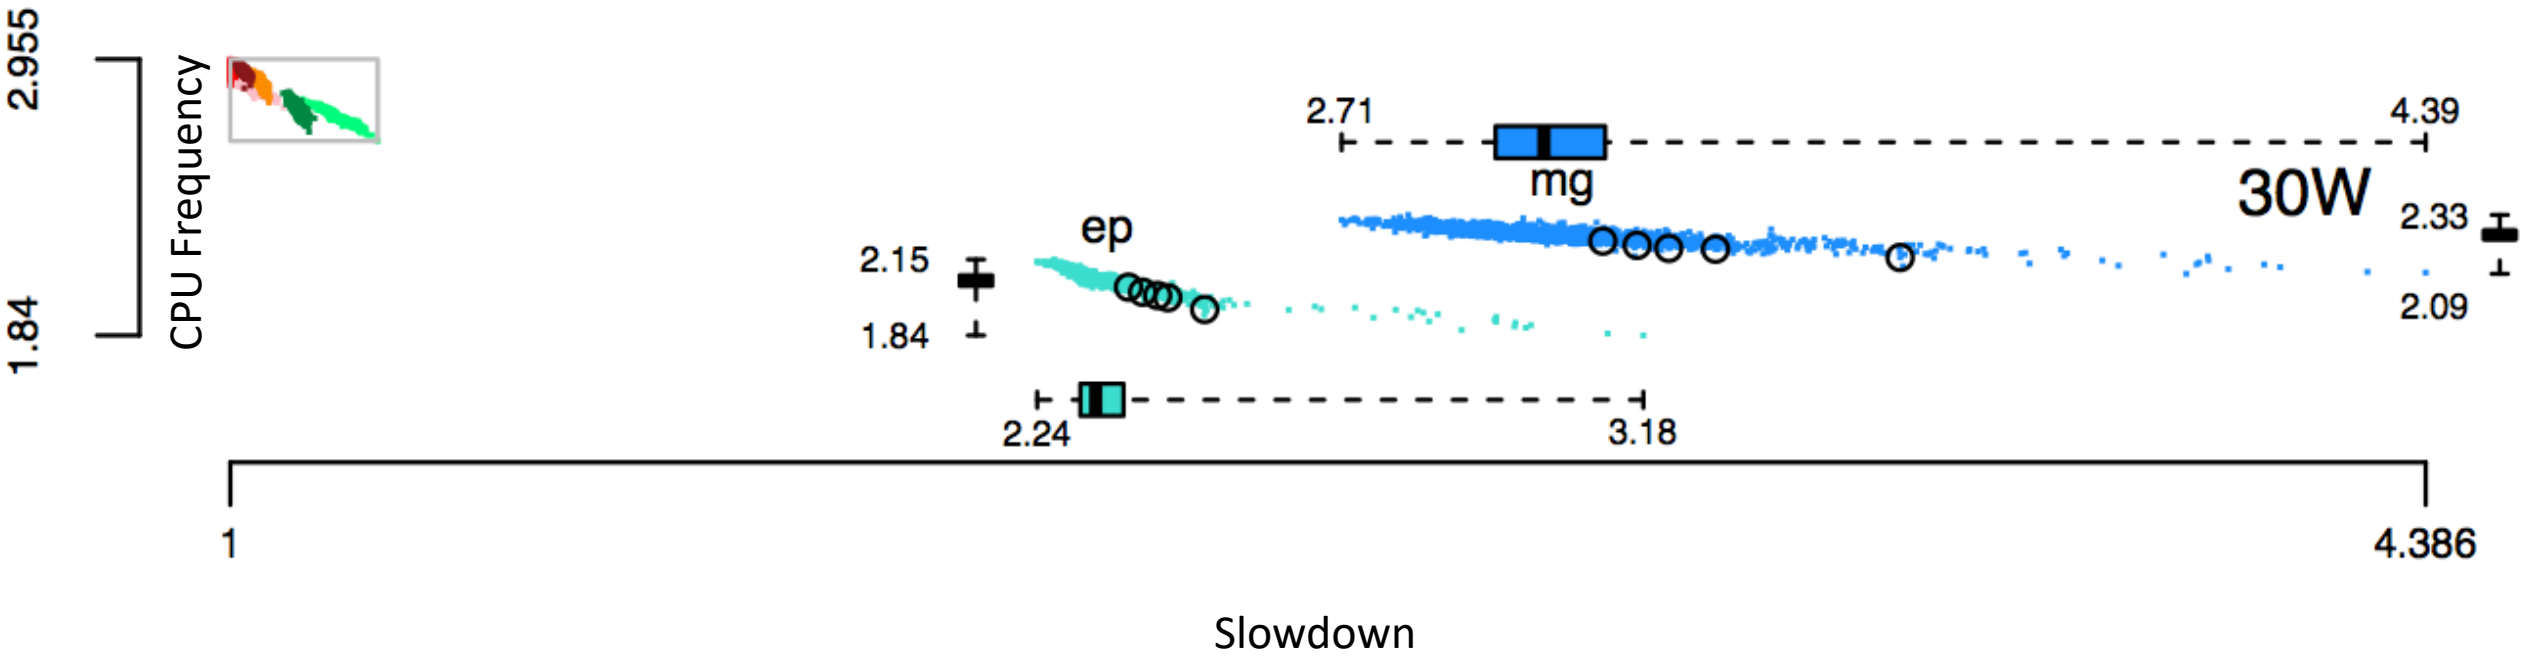

CASC

Center for Applied Scientific Computing

Lawrence Livermore National Laboratory

Disclaimer

This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.

TECHNISCHE UNIVERSITÄT DRESDEN

DRESDEN concept

Thomas Ilsche
Center for Information Services and High Performance Computing (ZIH)

Power and Performance Variations

Tenth Annual Workshop for the Energy Efficient HPC Working Group
EE HPC WG

ZIH
Center for Information Services & High Performance Computing

Comparison of Variations Across Different Processor Generations

- HPL (Linpack)
- Energy consumption for core phase
- Median of 10 repetitions

[Schuchart, J.; Hackenberg, D.; Schöne, R.; Ilsche, T.; Nagappan, R. & Patterson, M. K.
The shift from processor power consumption to performance variations: fundamental implications at scale
Computer Science - Research and Development, Springer Nature, 2016 , 31 , 197-205]

	Sandy Bridge-EP	Haswell-EP
Processor	E5-2690	E5-2680 v3
Sockets	2	
Nominal freq.	2.9 GHz	2.5 GHz
Turbo freq.	up to 3.8 GHz	up to 3.3 GHz
TDP	135 W	120 W
Memory	32 GB	64 GB
Sockets under test	412	1144
AVX frequency	--	2.1 GHz
Power measurement	Per-socket RAPL and per-node calibrated (HDEEM)	

Overview

HPL performance | Node power | Socket power

Sandy Bridge-EP (nominal 2.9 GHz)

Haswell-EP (nominal 2.5 GHz)

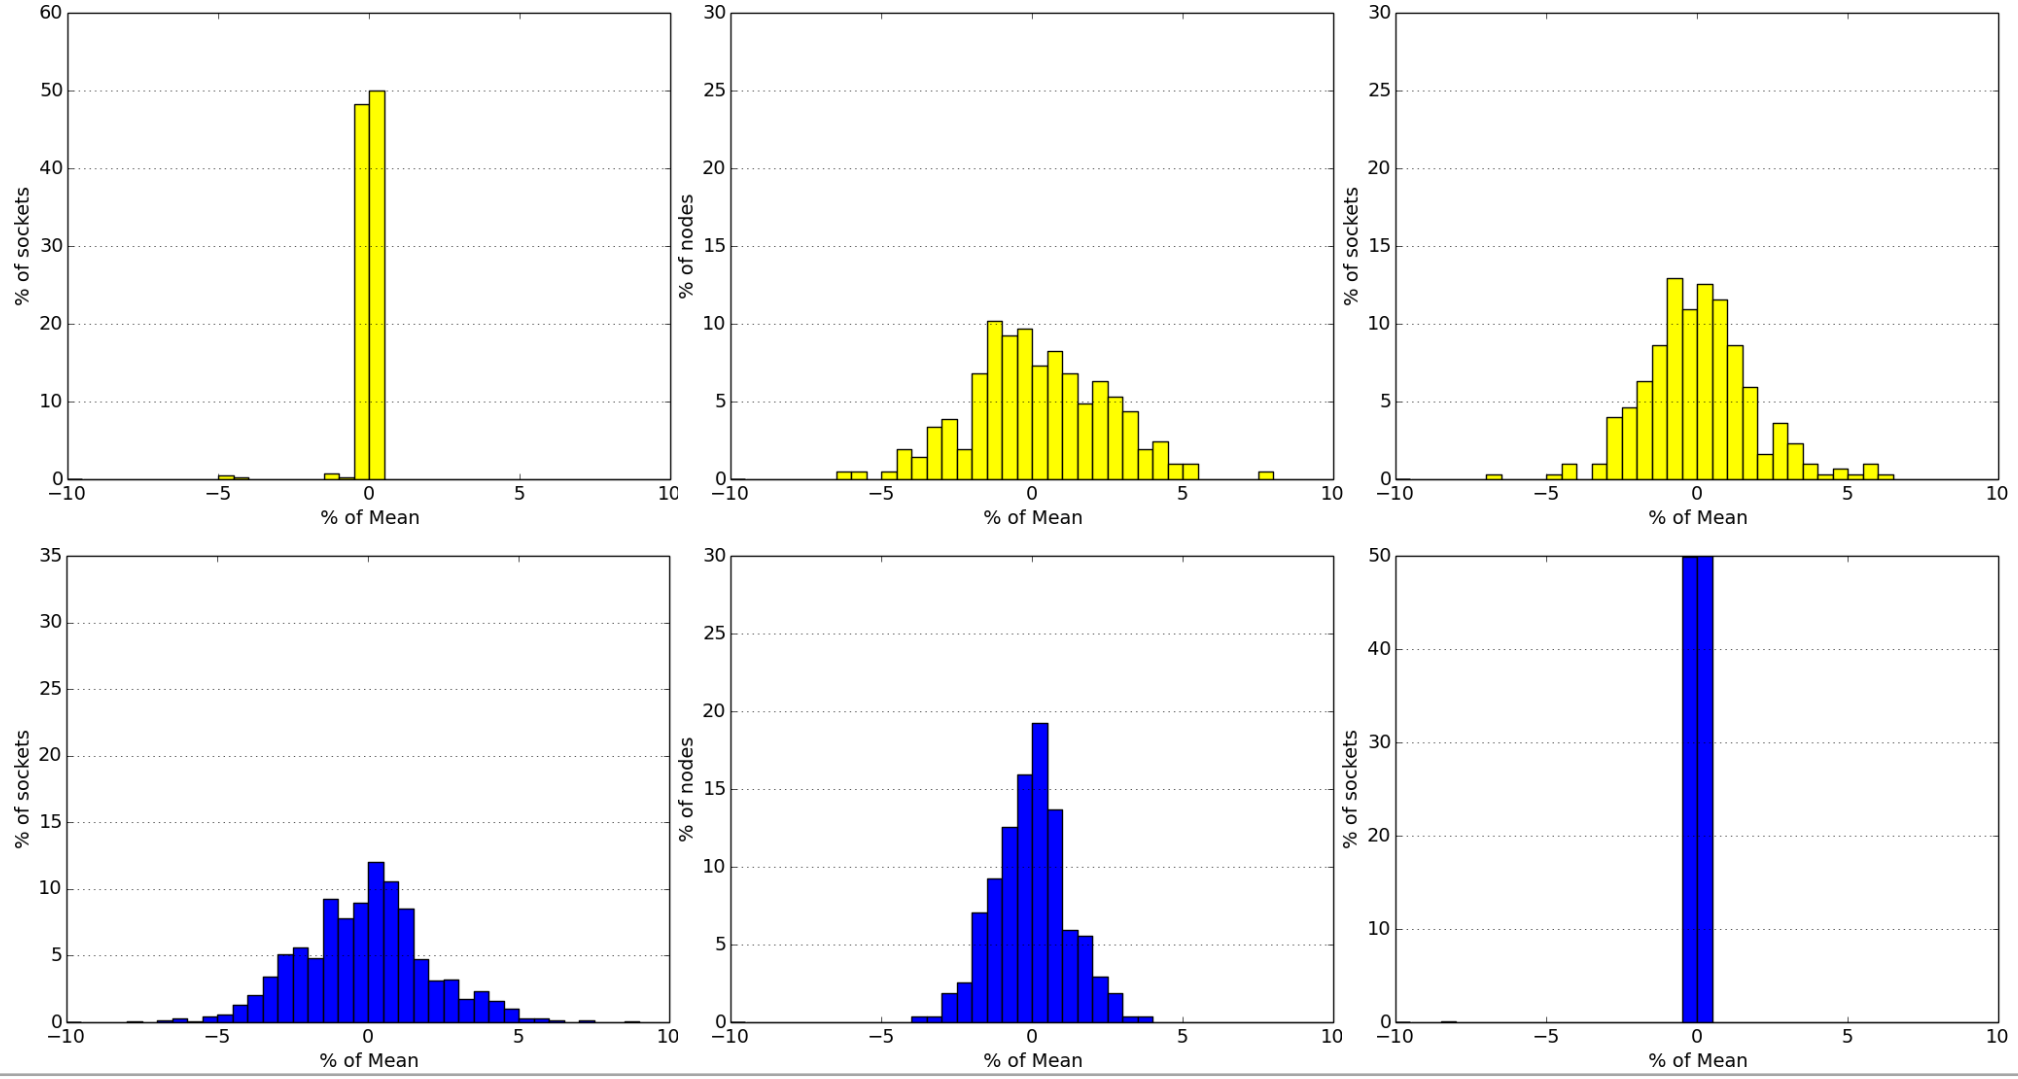

Power/Performance on Sandy Bridge-EP

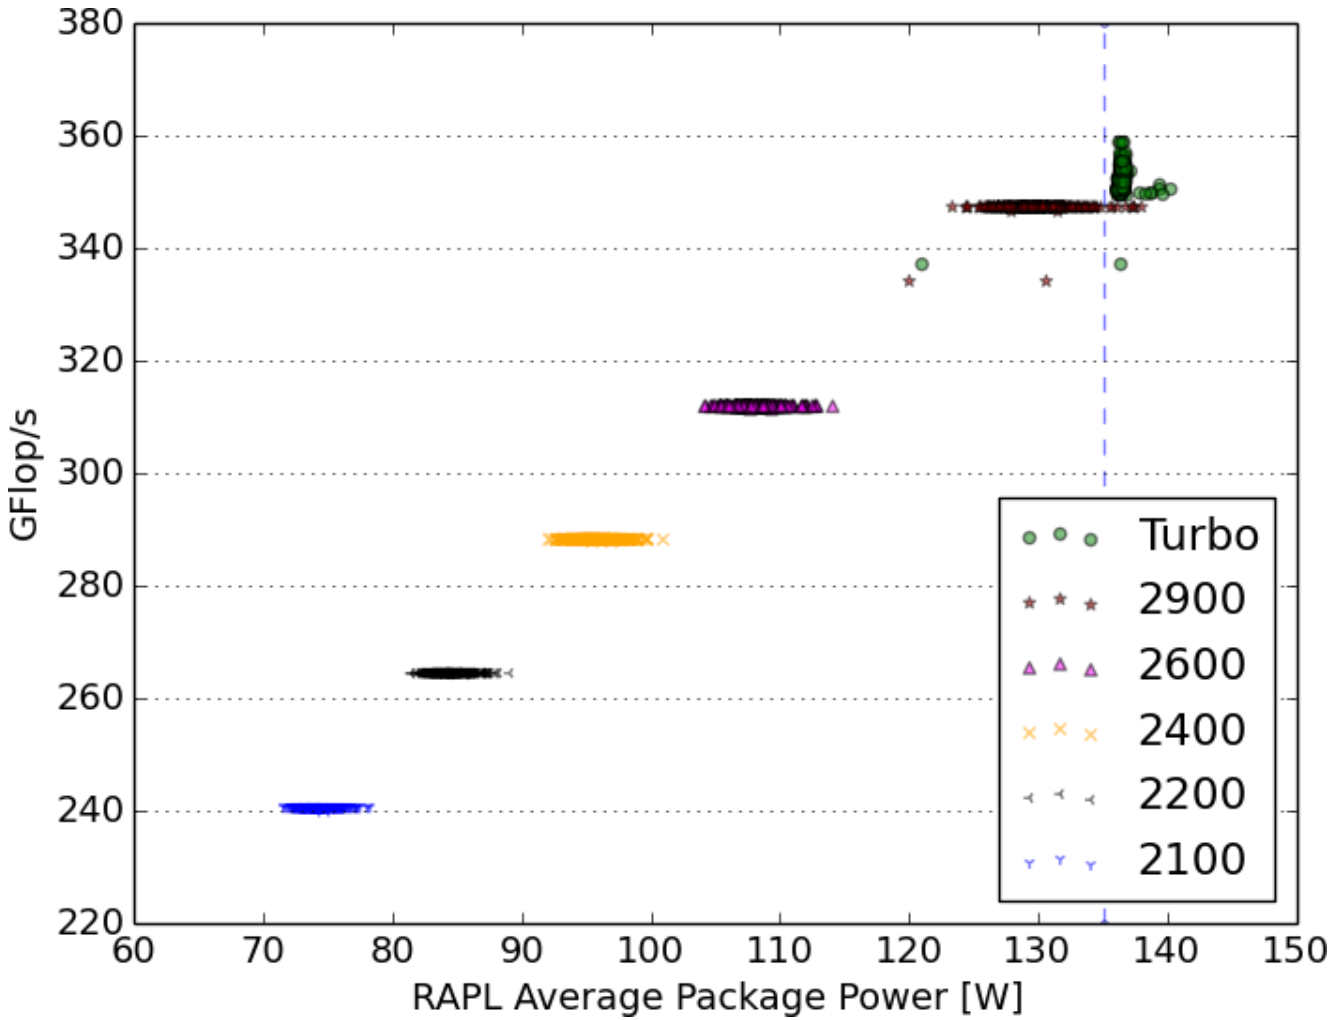

Power/Performance on Haswell-EP

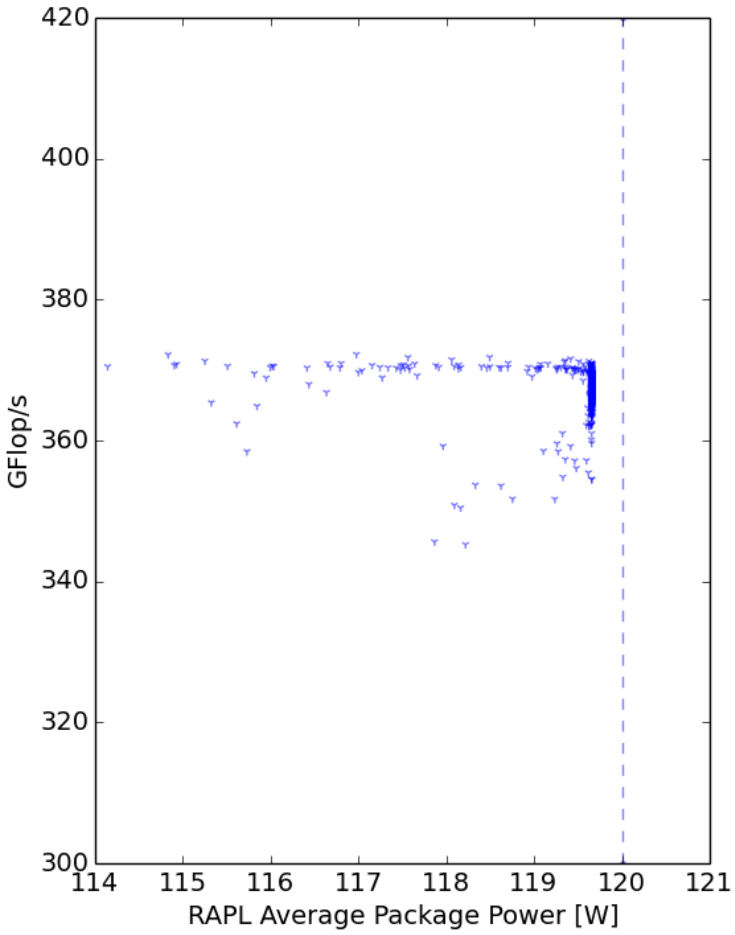

2.1 GHz

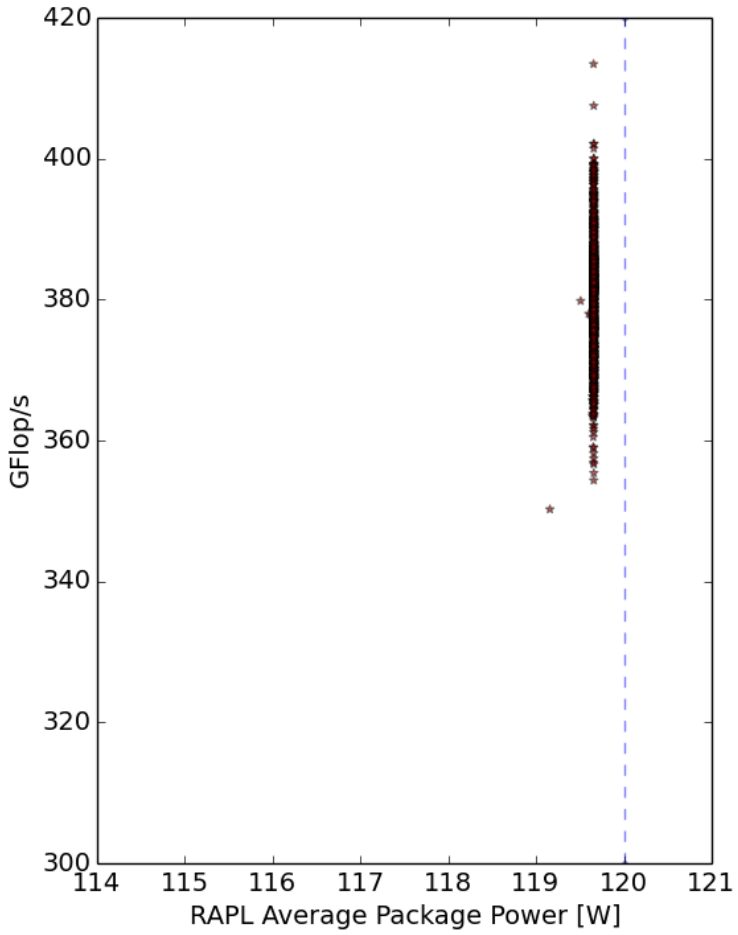

2.5 GHz

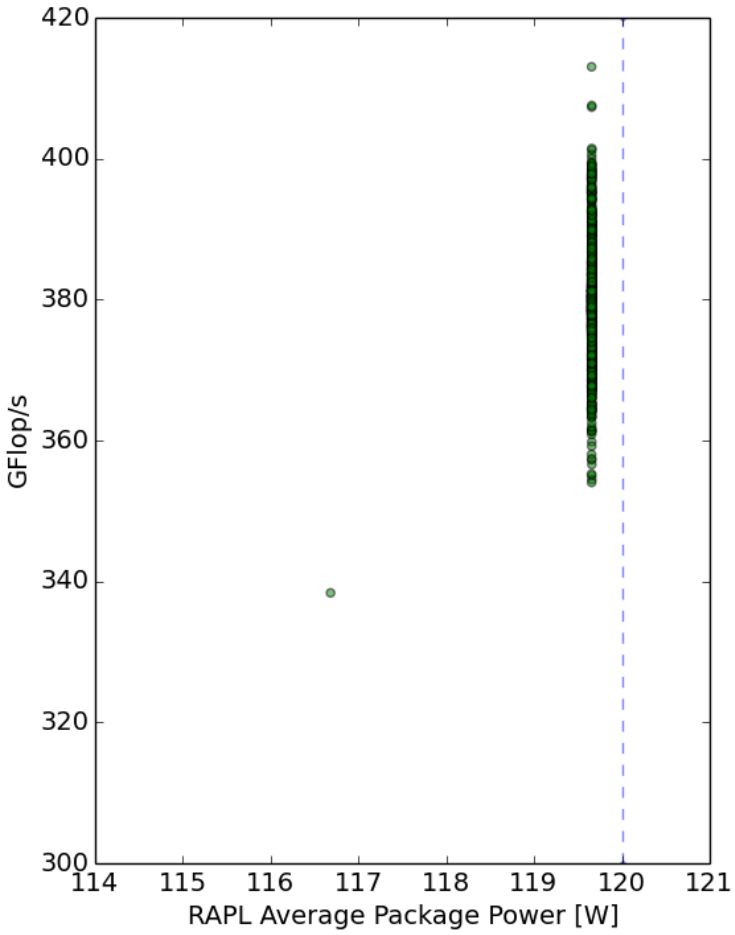

Turbo

The Shift from Power to Performance Variations: Implications

- Reproducible benchmarking
- Accurate power measurement for power capping
- Mitigation for synchronized parallel applications

The Shift from Power to Performance Variations: Questions

- Does this trend continue?
- Where will it end?
- Who mitigates the issue?
  - Applications – use it for load balancing
  - Site operators – hand out similar nodes through batch system
  - Vendors – more fine-grained power-focused binning

Bonus Variation: Data-dependent Power Consumption

The same workload for four minutes – 58 W difference

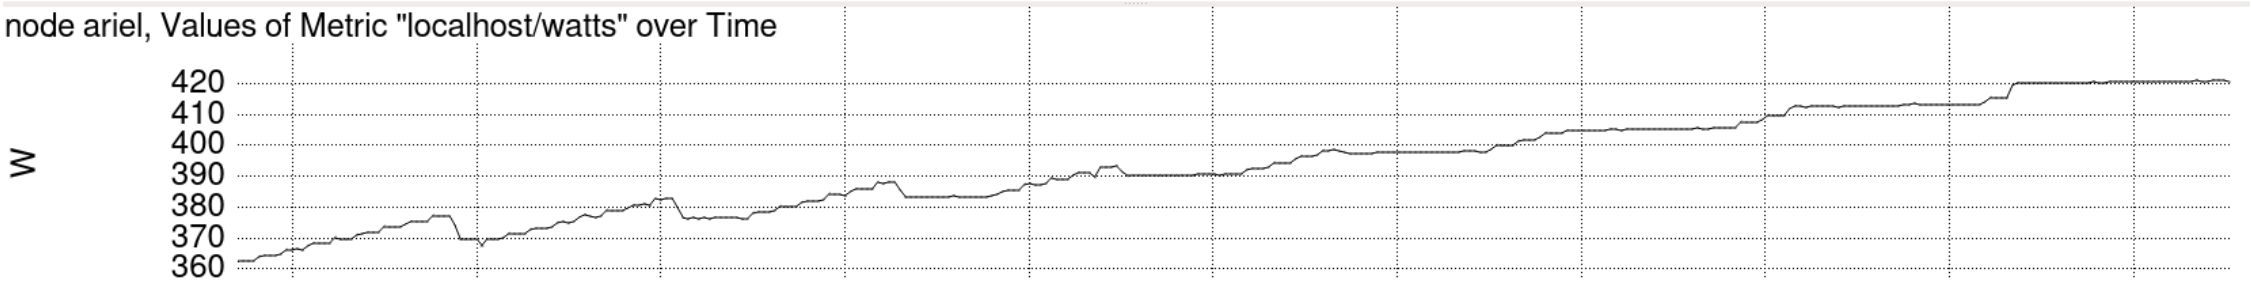


```
vxorps %zmm0,%zmm8,%zmm8
vxorps %zmm1,%zmm9,%zmm9
…
vxorps %zmm7,%zmm15,%zmm15
```

SKL-SP @ 3 GHz

[Schöne, R.; Ilsche, T.; Bielert, M.; Gocht, A. & Hackenberg, D.
Energy Efficiency Features of the Intel Skylake-SP Processor and Their Impact on Performance
accepted for publication]

Even Stronger Impact Possible

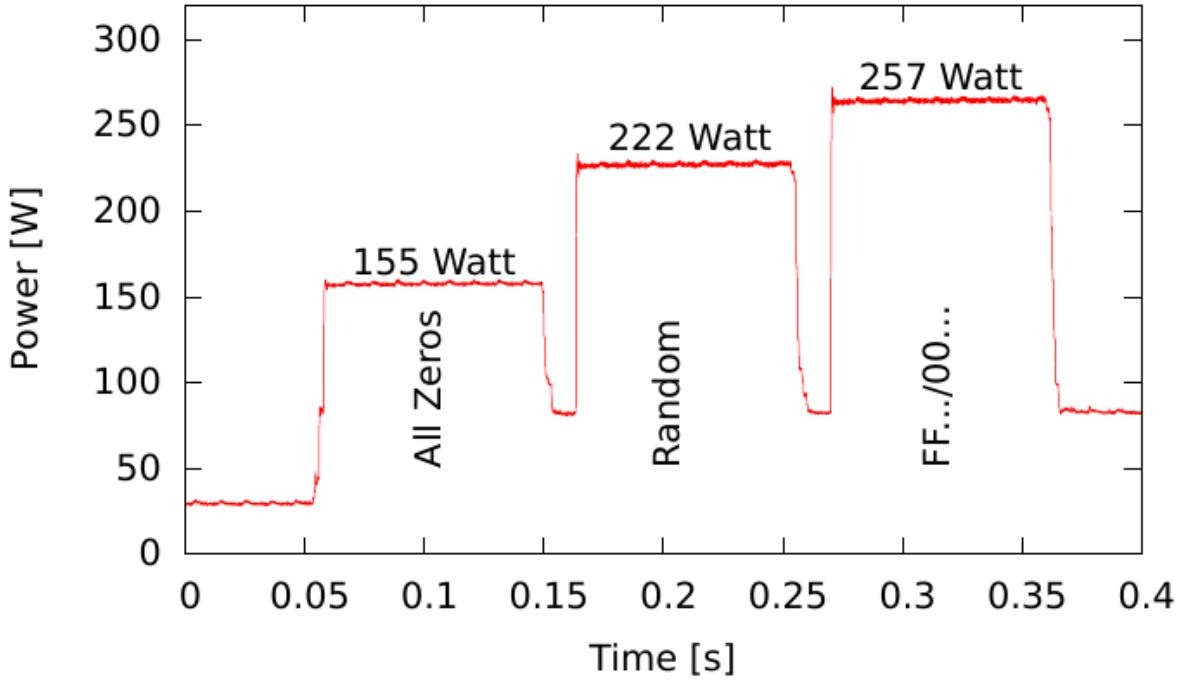

Fig. 1. Power Consumption of FMUL test kernel with three different input vectors on GTX580

[J. Lucas and B. Juurlink, ***ALUPower: Data Dependent Power Consumption in GPUs***,
2016 IEEE 24th International Symposium on Modeling, Analysis and Simulation of Computer and Telecommunication Systems (MASCOTS), London, 2016, pp. 95-104.
doi: 10.1109/MASCOTS.2016.21]

Bonus Variation: Data-dependent Power Consumption

- High core count
- Wide SIMD instructions
- Power increases with number of set bits

→ Is this something we need to care about?

→ Impact on energy modeling!

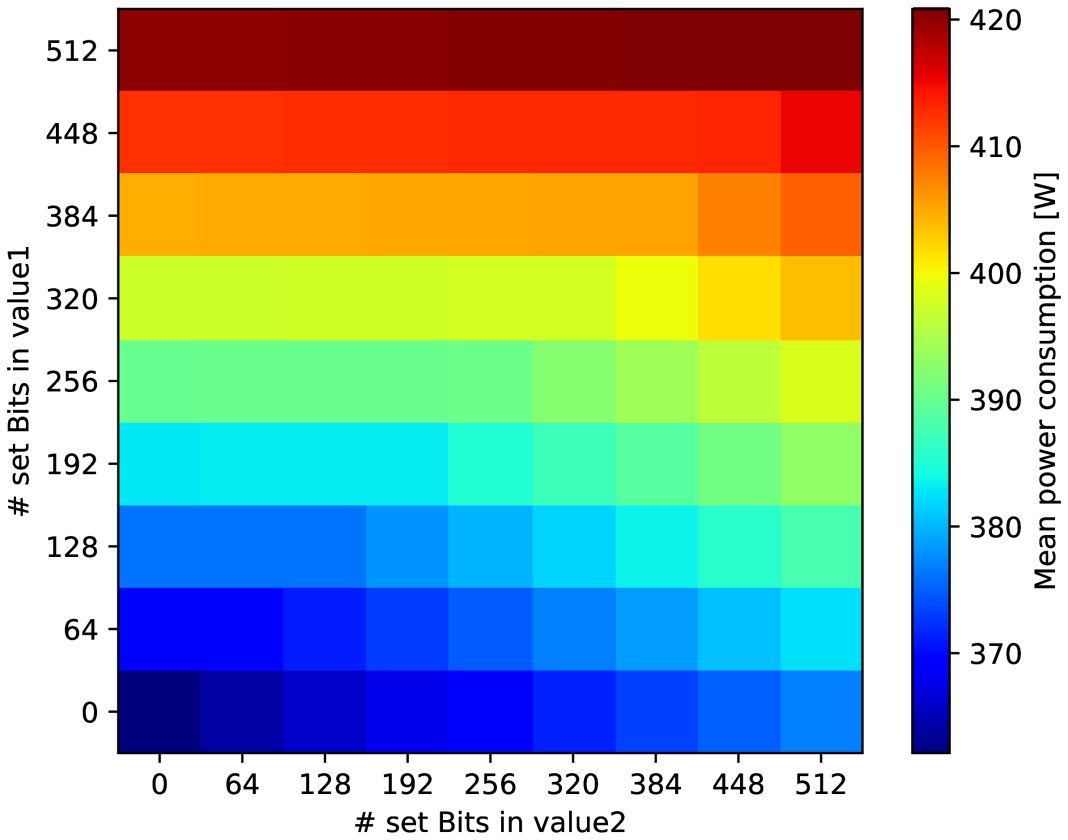

Q&A

Challenges do to Manufacturing Variation