Benchmarking For Processor Performance Variations: A Case Study With An Arm Processor

James (Xinhua) Lin

Vice Director of HPC Center

Shanghai Jiao Tong University

BID Workshop @ PPoPP 2022

Shanghai Jiao Tong University

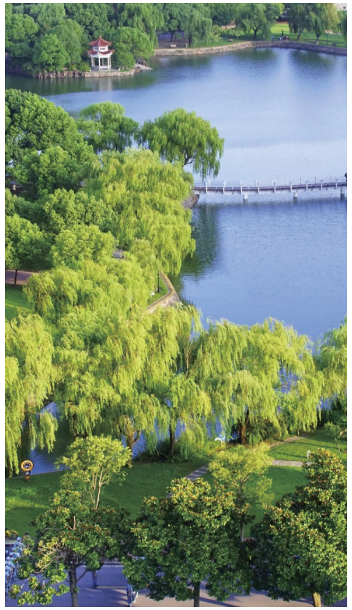

- Establishment in 1896, Shanghai Jiao Tong University (SJTU) is 2nd oldest and one of TOP5 universities in China.
- The university has hosted CCGrid2009 and IPDPS2012, both are the first time in China.

Siyuan Mark-1, the Fastest in China Universities

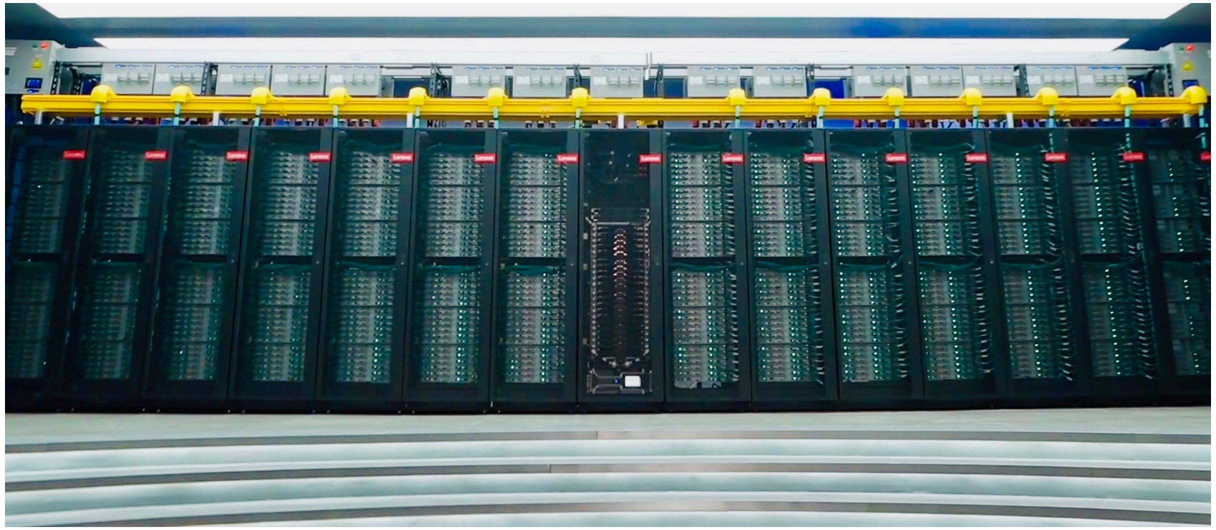

Donated by Lenovo CEO, an alumni of SJTU, in 2021

上海交通大學
SHANGHAI JIAO TONG UNIVERSITY

Outline

The hardware configuration of π-2.0

of Nodes: 656 blade nodes.

CPU: Intel Xeon-6248（20 cores） x 2

Memory: 16GB DDR4 2666MHz x 12

Network: 100Gbps, Intel Omni Path

OS: CentOS 7.7 (3.10.0-1062）

Online since Oct. 2018.

https://hpc.sjtu.edu.cn/Item/Hardware.htm

The performance variation on π-2.0

The single-node HPL performance varies among nodes, and the performance gap between the worst and the best performance is over 200 GFLOPS.
Over 17% single-node performance gap compared to the best performance.

<2075 GFLOPS

2290GFLOPS

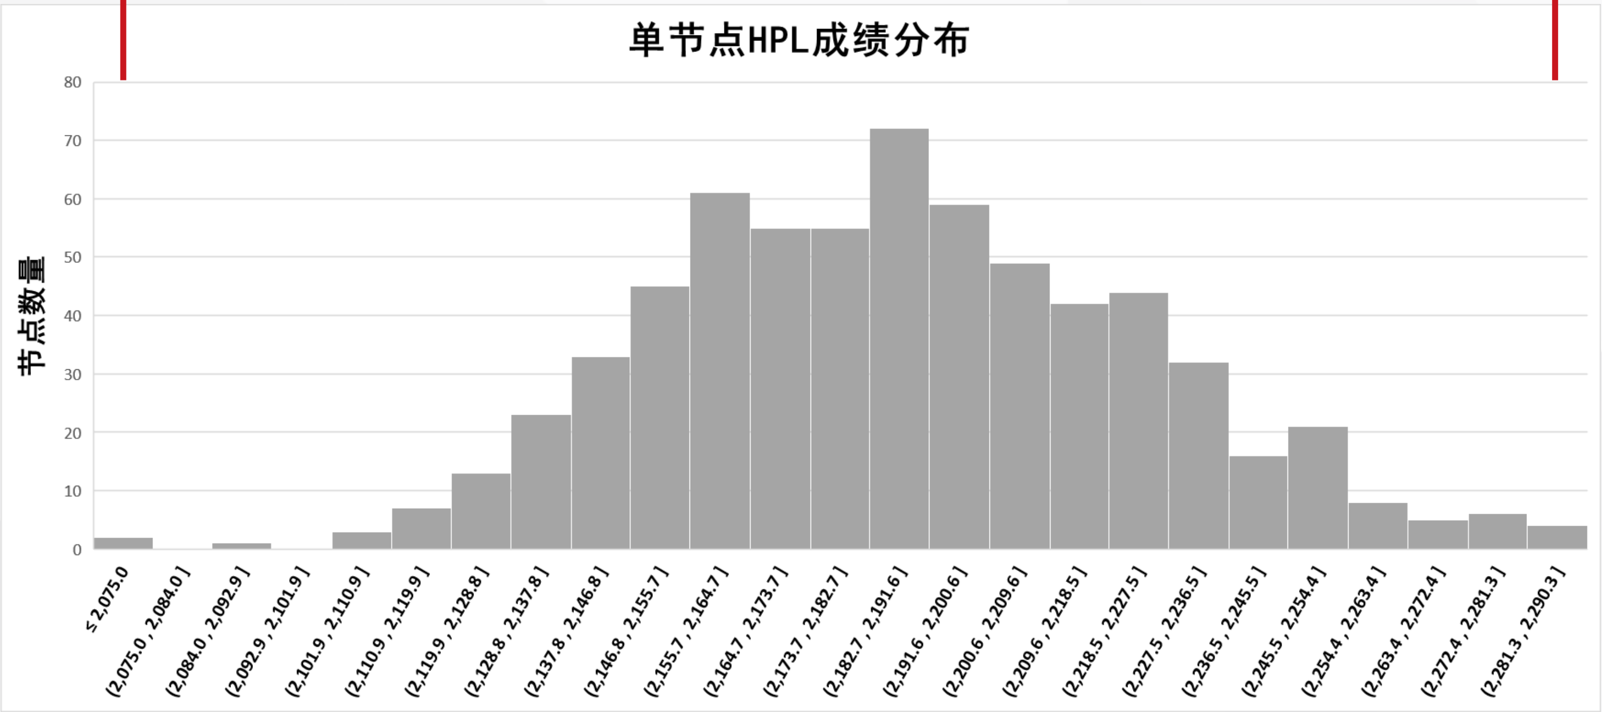

The performance variation on π-2.0: what we did

By further benchmarking, we found that the performance variation is related to the variation in CPU frequency, and memory bandwidth.

Conclusions：

1. Several CPUs 'bad quality leads to the frequency drops.
2. Several nodes have 2-rand and 1-rand DDR4 memory mixed.
3. The performance anomaly has been fixed after replacing all problematic components.

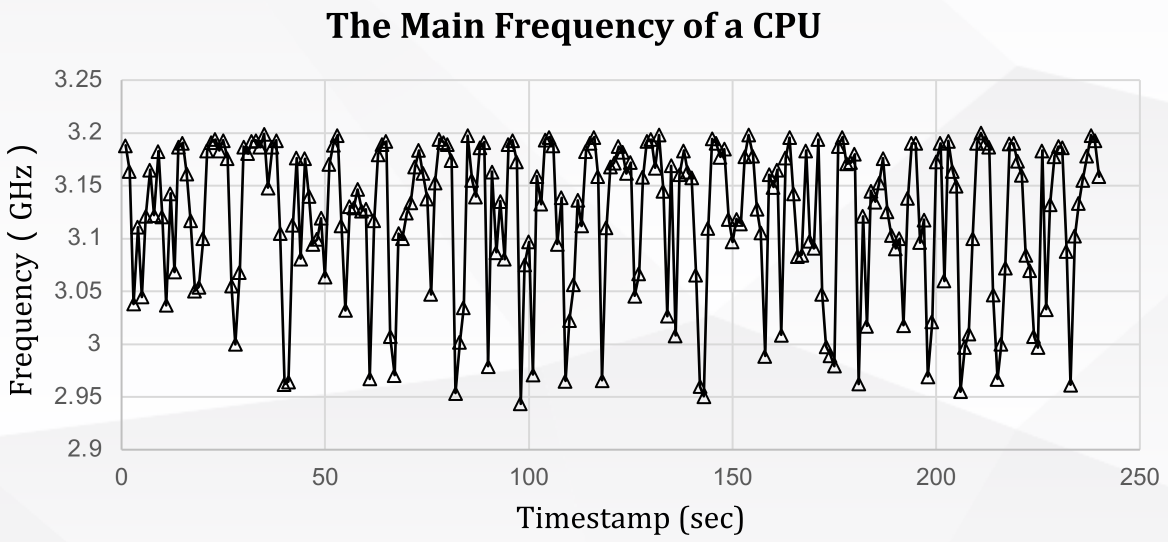

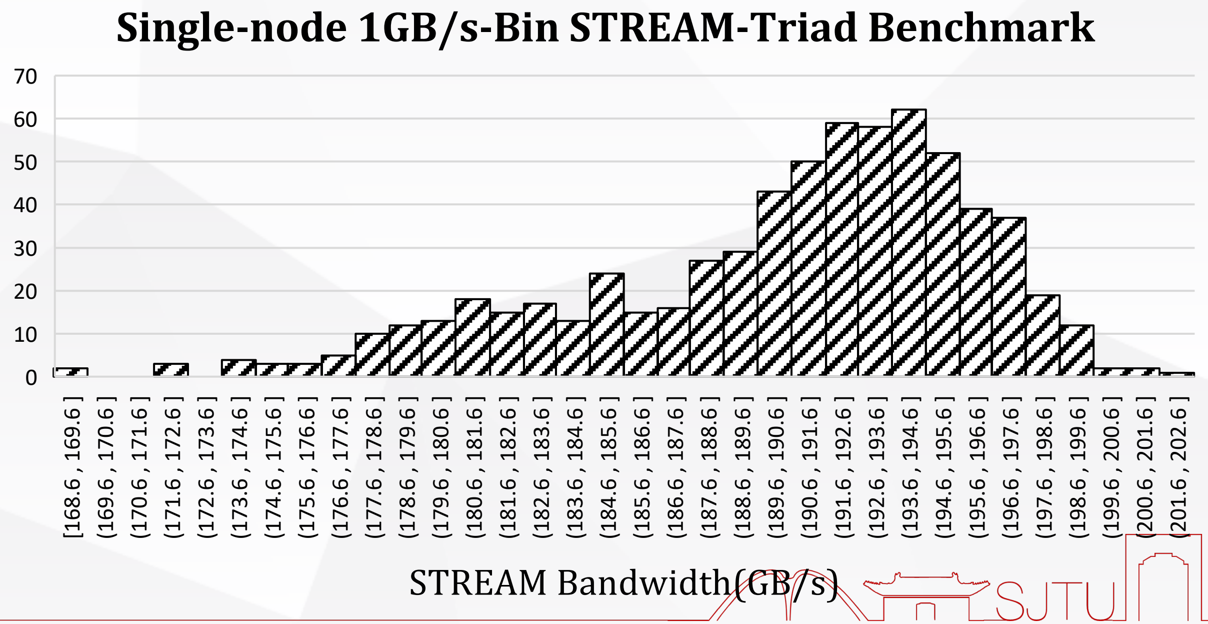

SJTU

The covert harm of performance variations

Performance variations are often led by the defect in the software, the hardware, and the algorithm on HPC systems. The sources of performance variations widely exist in different hierarchies of HPC systems. And the performance degradation is difficult to predict and measure after the coupling of performance variation sources.

The classification of hierarchical phenomenon of performance variations.

- Cluster-level: The performance varies in the whole cluster or several cabinets.
  - e.g.: The running time of sequenced similar jobs varies during weeks.
- Node-level: Inter-node performance variations.
  - e.g.: Node-to-node performance variation in BSP applications.
- OS-level: The program is interrupted by the operating system.
  - e.g.: Unstable CPU utilization due to OS noise.
- Chip-level: Flaw in microarchitecture leads to unstable inter-loop performance。
  - e.g.：Inconsistent performance in a for-loop.

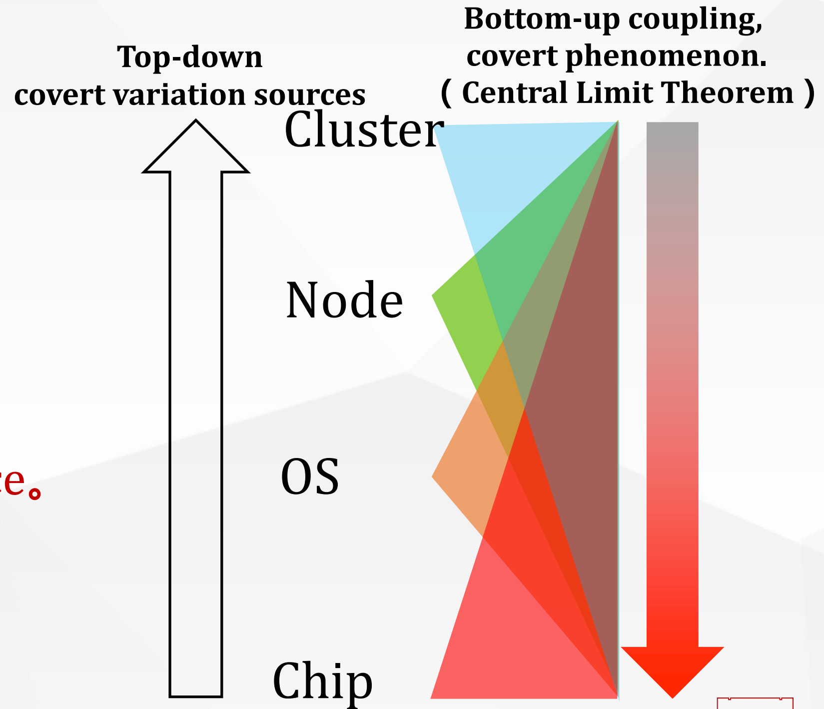

Related work 1: Chip-level performance variation

HPC Systems: Stampede-2 （Top500 #44, 2021/11）

John McCalpin

HPL and DGEMM performance variability on the Xeon Platinum 8160 processor (SC2020)

The phenomenon:

Several random tests slow down during the repeating single-node HPL benchmark.

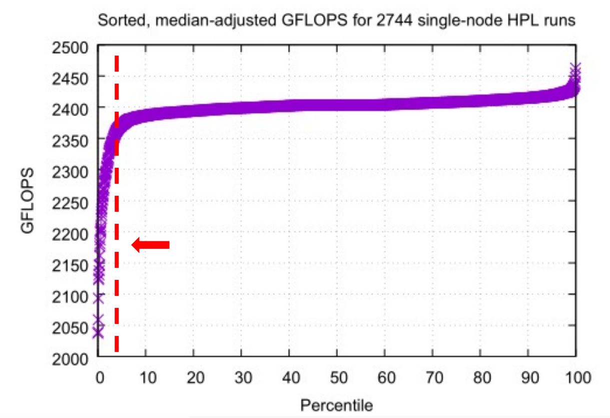

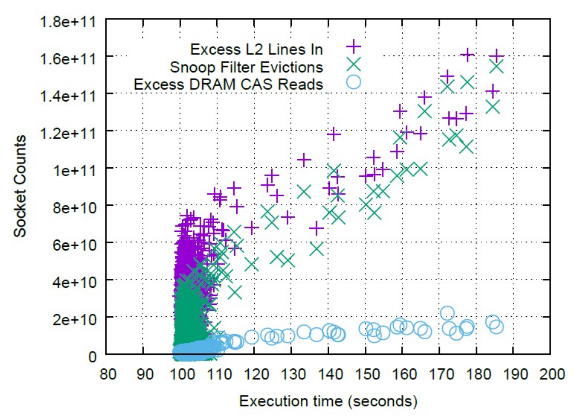

The sources:

The flaw in the design of the Coherence Agent (CHA) of Intel Platinum 8160 leads to occasional cache line false evict.

06/2018	15	PowerEdge C6320P/C6420, Intel Xeon Phi 7250 68C 1.4GHz/Platinum 8160, Intel Omni-Path	DELL EMC	367,024	10,680.7	18,309.2
11/2017	12	PowerEdge C6320P/C6420, Intel Xeon Phi 7250 68C 1.4GHz/Platinum 8160, Intel Omni-Path	DELL EMC	368,928	8,317.7	18,215.8

Related work 2: OS-level performance variation

HPC Systems: Jaguar(Top500 #2, 2009/06)
Torsten Hoefler et. al.
Characterizing the Influence of System Noise on Large-Scale Applications by Simulation (SC2010)

The phenomenon:
The performance degradation of MPI collective operation keeps worsening as the applications on Jaguar scale up.

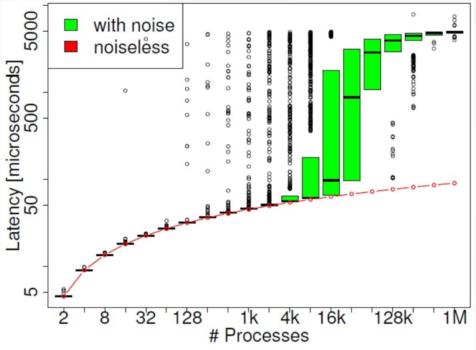

(e) 10x Jaguar

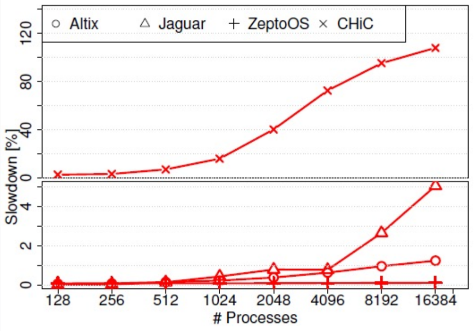

(a) POP

The sources:
Interrupts from the operating system jeopardize the inter-node synchronization, then lead to occasional slowdown on the node.

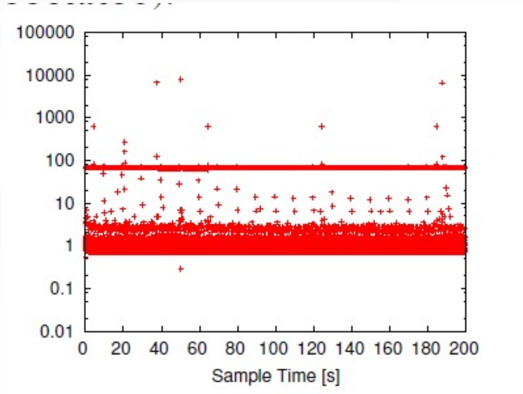

(d) Jaguar

Related work 3: Node-level performance variation

HPC Systems: Cab(Top500 #391, 2016/06), **Edison**(#292, 2020/06), **Stampede**(#20, 2017/06)
Bilge Acun, et. al.
Variation Among Processors Under Turbo Boost in HPC Systems (SC2010)

The phenomenon:
The single-node performance of MKL-DGEMM varies in all three HPC.

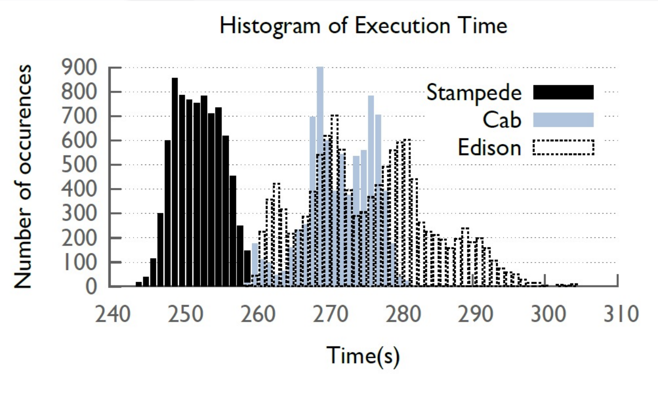

The sources:
MKL-DGEMM has high utilization of 512-bit ALUs, and triggers frequent throttling. The different quality and working environment of processors leads to different throttling levels, which induce the node-level performance variation.

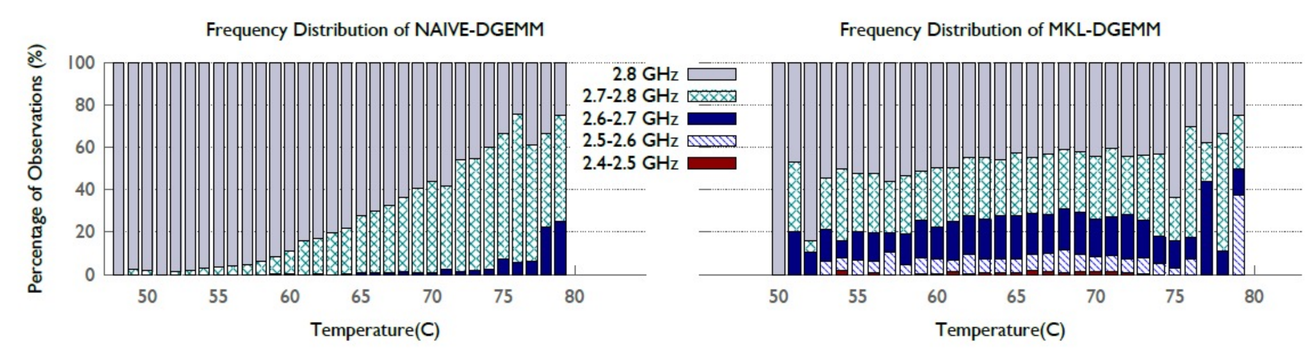

Related work 4: Cluster-level performance variation

HPC System：Cori(Top500 #37, 2021/11)
Abhinav Bhatele et. al.
The Case of Performance Variability on Dragonfly-based Systems (IPDPS2020)

The phenomenon:
During the 5-month data collection on Cori, the performance variation in 4 applications is detected, with up to a 300% performance gap.

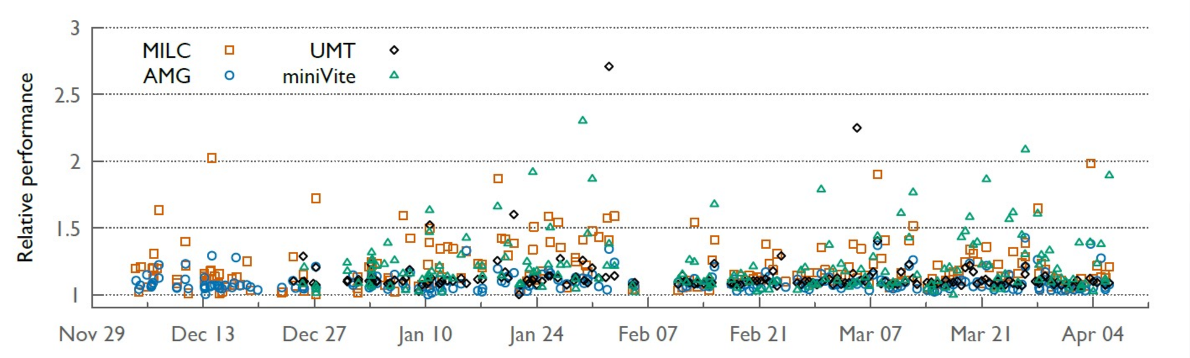

The sources:
The study locates the performance variation source which is the coupling of the dragonfly network structure and the job scheduling strategy. An inappropriate node schedule brings around a 200-second slowdown in the communication of the 512-node AMG job.

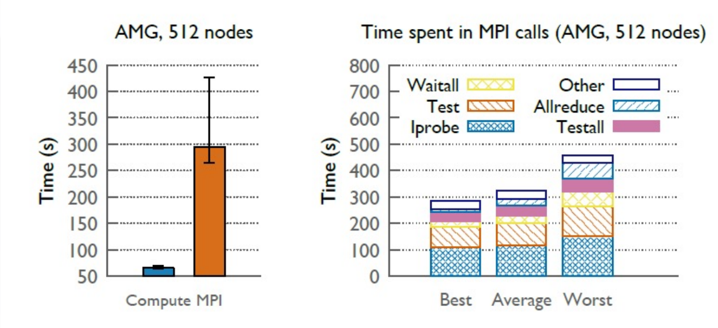

上海交通大學
SHANGHAI JIAO TONG UNIVERSITY

Outline

1. Background
2. The performance variation of KP920
3. Case study: SJTU KP920 HPC system
4. Conclusion

The experimental platform

Platform	Xeon6148	KP920
CPU	Intel Xeon Gold 6148	HiSilicon Kunpeng 920
# of Cores	20	64 (Some data from 48-core SKU)
Frequency(GHz)	2.2 (AVX512)	2.6
# of Socket	2	2
Memory Size(GB)	768（12 x 64 GiB）	256 (16 x 16 GiB)
Memory frequency(MHz)	2666	2666
OS Version	CentOS-7.3 Kernel 3.10.0	CentOS-7.3 Kernel 4.18.0
Compiler	GCC-8.2.0	GCC-8.2.0
MPI Library	MVAPICH2-2.3	MVAPICH2-2.3
BLAS Library	OpenBLAS-0.3.4	OpenBLAS-0.3.4

The top-down benchmark for diagnosing performance variation

Mini Application

HPL SNAP TeaLeaf CloverLeaf

⇩

µArch Benchmark

STREAM BenchIT LIKWID OSU Benchmark

⇩

µArch Monitoring

PAPI LIKWID

Step 1：Mini-App benchmark
Identifying performance anomalies with mini-apps that have special performance characteristics.

Step 2：µArch benchmark
Further benchmarking with micro benchmarks.

Step 3: Analysis of PMU counters
Using PMU counters to build correlation for further diagnosing.

The variation in the CPI of arithmetic instructions

Evaluation:

Evaluating the performance of ALUs with basic assembly arithmetic instructions.

Phenomenon：

1. The throughput of FP64 FMA and MUL instruction is halved.
2. The FP16 DIV and SQRT instructions are not hard-wired implemented
3. As a socket-to-socket comparison, the FP32 arithmetic on KP920 is better than Xeon6148.

Conclusion:

1. Performance of mix-precision applications will vary in different implementations.
2. Arithmetics with different precision will meet inconsistent performance.

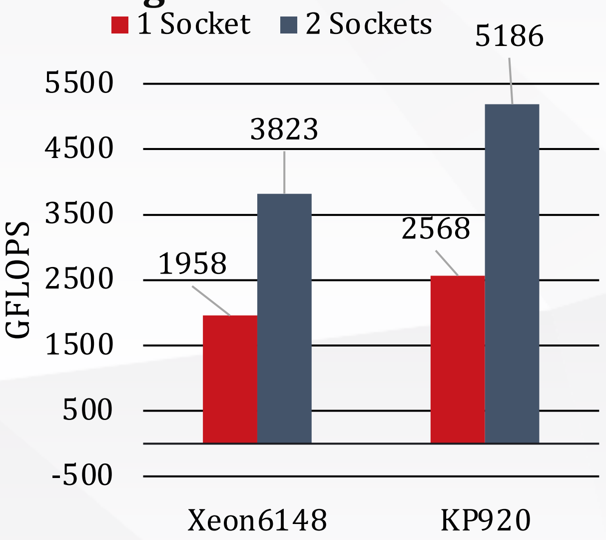

The CPI of basic arithmetic instructions

Instruction	FP16		FP32		FP64	
	Absolute	Pipelined	Absolute	Pipelined	Absolute	Pipelined
ADD	2	0.5	2	0.5	2	0.5
SUB	2	0.5	2	0.5	2	0.5
MUL	4	0.5	5	0.5	5	**1**
DIV	**46**	**42**	16	12	16	12
FMA	4	0.5	5	0.5	6	**1**
SQRT	46	42.15	18	14.37	18	15.55

SNAP: The multi-layer gap in the latency of the memory subsystem

Evaluation:

1. Evaluating the memory subsystem with SNAP, utilizing all cores of a single node.
2. Two datasets. A small dataset(c4096) with better data locality, and a big dataset(c32768) with more cache miss and refill.
3. On KP920, observe the change of the performance along with the process/thread ratio.

Phenomenon：

1. On KP920, the performance drops when running the larger dataset.
2. The best performance appears when running 4 threads with each process.

Further micro benchmarking:

1. Evaluating the memory bandwidth and cache latency for further diagnosing the performance drop in the large dataset.
2. Evaluating the latency of inter-core data sharing.

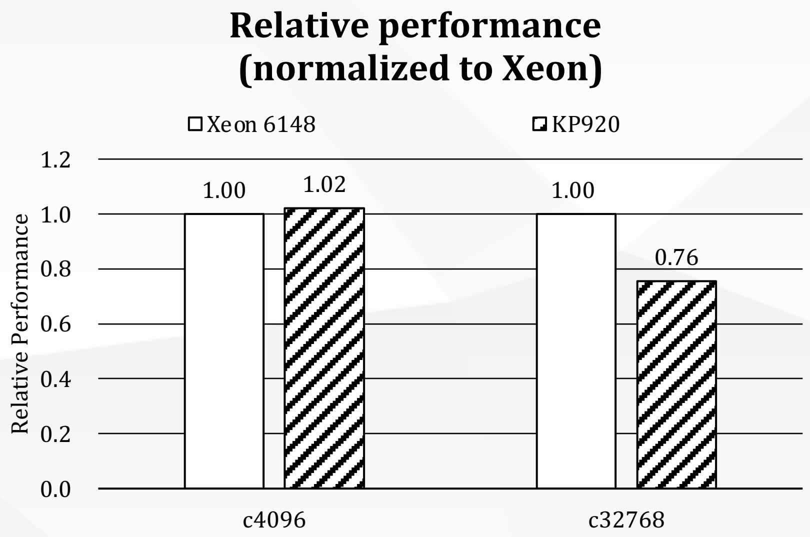

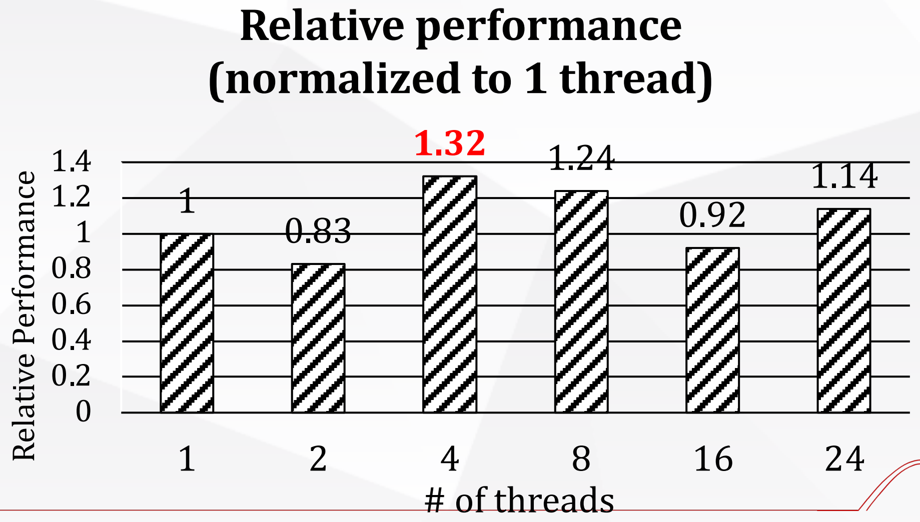

SNAP: The multi-layer gap in the latency of the memory subsystem

A further benchmark for Ph.1:

1. Evaluating the memory bandwidth and cache latency for further diagnosing the performance drop in the large dataset.

Conclusion：

1. On KP920, the memory bandwidth drops in STREAM-add and STREAM-triad kernel, which implies that the compute kernel with high arithmetic intensity is difficult to utilize the DRAM bandwidth.
2. The unstable L3 cache latency and the larger DRAM latency jeopardize the performance when data access of the sweep process of SNAP hits the L3 and DRAM memory in the larger dataset.

Dual-socket STREAM_MPI benchmark

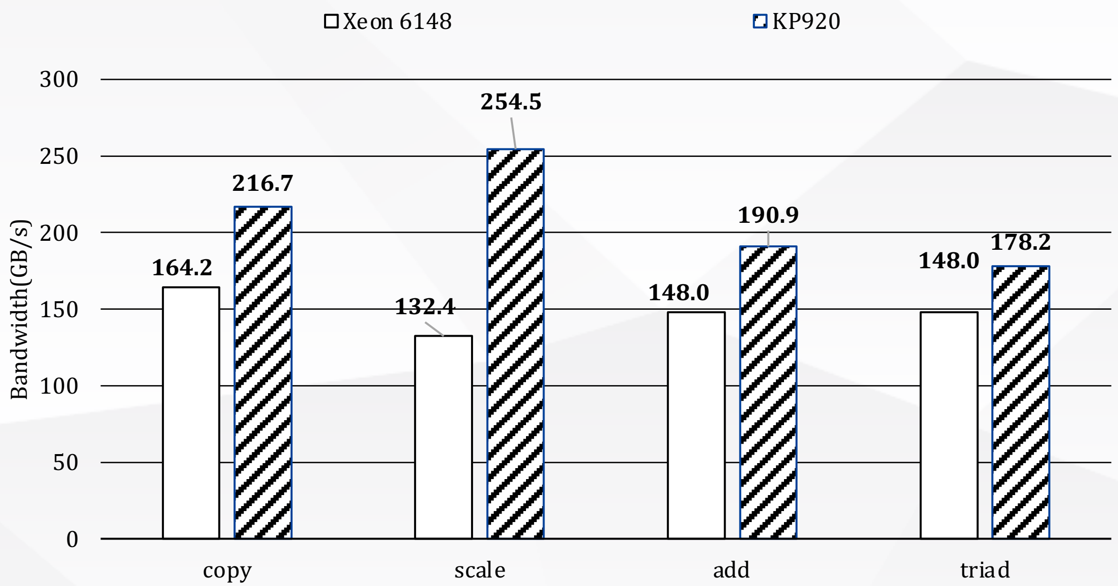

Cache Latency (ns)

Chip	L1	L2	L3	DRAM
Xeon6148	1.295	4.532	22.83	88.611
KP920	2.000	4.002	19 ~ 44	125.00

SNAP: The multi-layer gap in the latency of the memory subsystem

A further benchmark for Ph.1 and 2:

1. Evaluating the latency of inter-core data sharing.

Conclusion：

1. By making use of the core affinity in the contiguous 4-core cluster (CCL), the latency of Intra-CCL access is notably lower than in the other scenarios.
2. When the data sharing goes across a die, the latency is notably higher than intra-die access. Besides, L2 cache inter-die access is notably slower than one in Xeon 6148 processor.

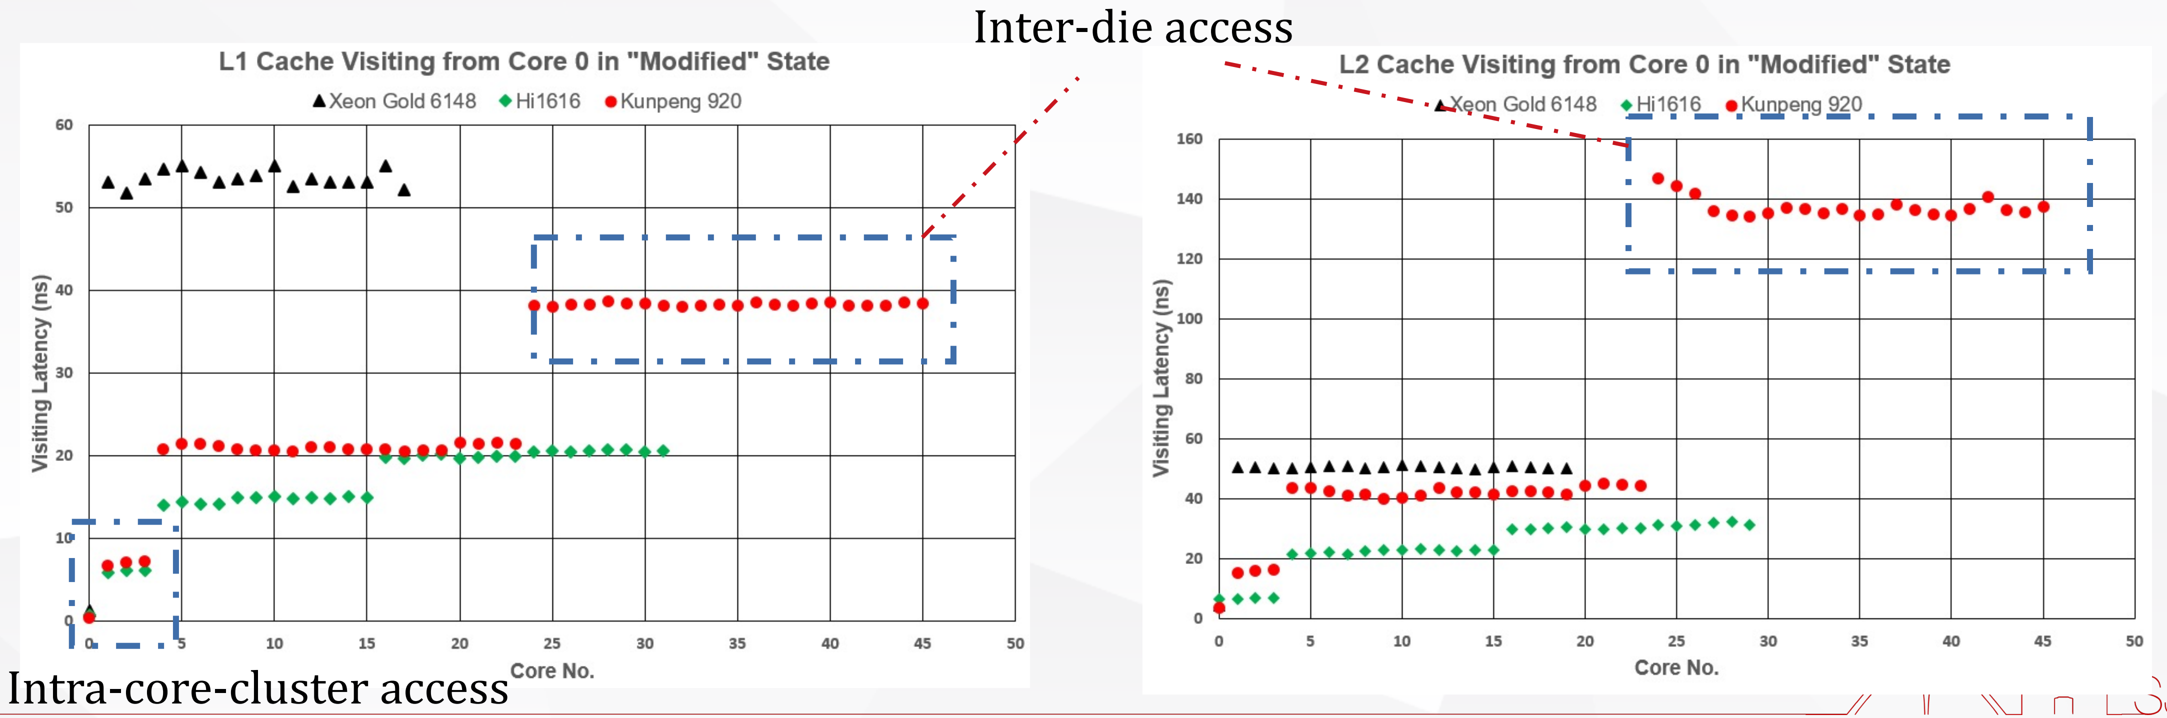

TeaLeaf: Run-to-run performance variation

Evaluation:

1. TeaLeaf is a memory-bound proxy application that computes 3 different stencil kernels.
2. Running each process with 4 threads mapping to 4 neighbor cores according to the experience in SNAP tests.

Phenomenon：

1. In 100 single-node tests, KP920 shows run-to-run performance variation with up to a 3.3% performance drop compared to the best performance. (1.2% on Xeon6148)

Further micro benchmarking:

1. The stencil kernels in the test case put stress on L3 cache sharing and DRAM memory.
2. Repeatedly benchmarking L3 inter-core access for further investigation.
3. Repeatedly benchmarking multi-core memory bandwidth.

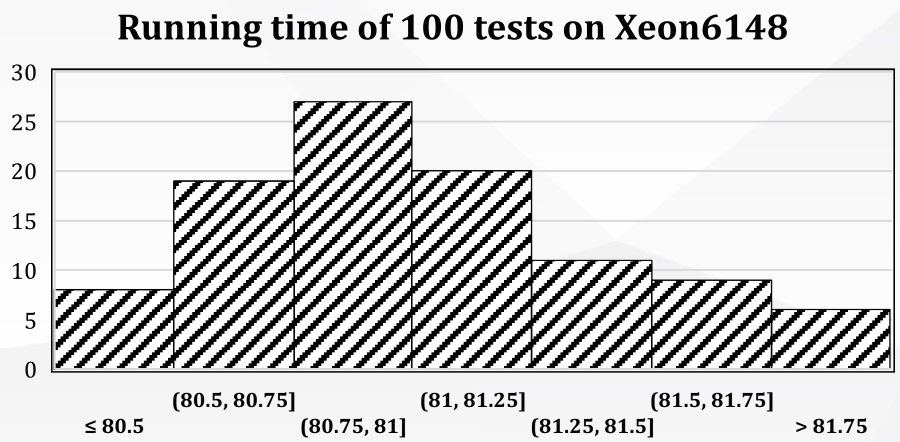

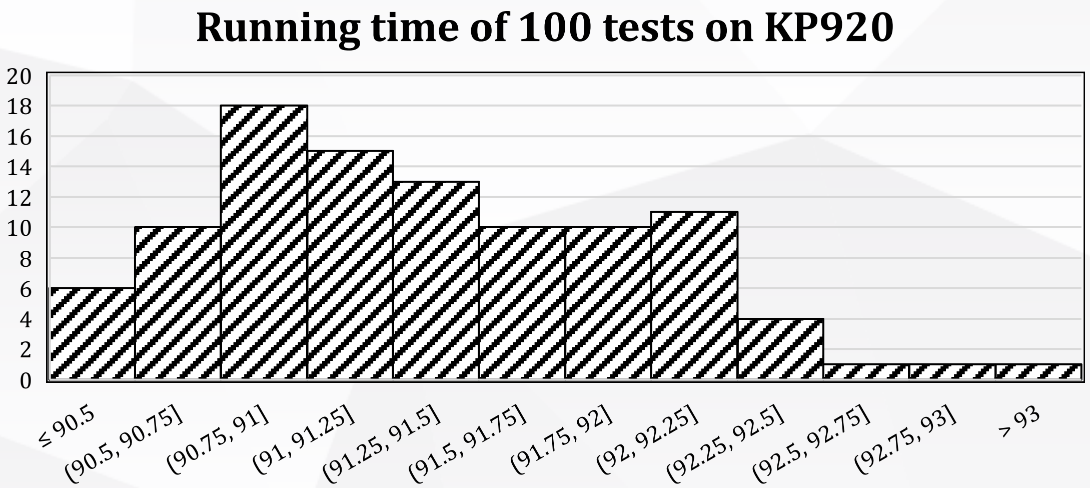

SJTU

TeaLeaf: Run-to-run performance variation

A further benchmark for Ph.1:

1. Repeatedly benchmarking L3 inter-core access for further investigation.

Phenomenon：

1. The latency of L3 cache sharing incurs notable variation.

Conclusion:

1. The latency of L3 cache sharing is a performance variation source on KP920.

L3 remote access on KP920

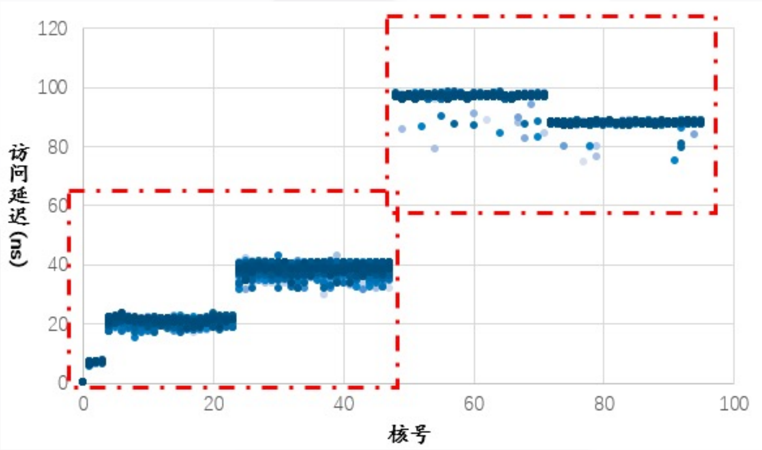

TeaLeaf: Run-to-run performance variation

A further benchmark for Ph.1:

Repeatedly benchmarking multi-core memory bandwidth. Recording the minimum, the average, and the best bandwidth in a 1- to all-core STREAM benchmark.

Phenomenon：

The real multi-core memory bandwidth on KP920 incurs run-to-run performance variation.

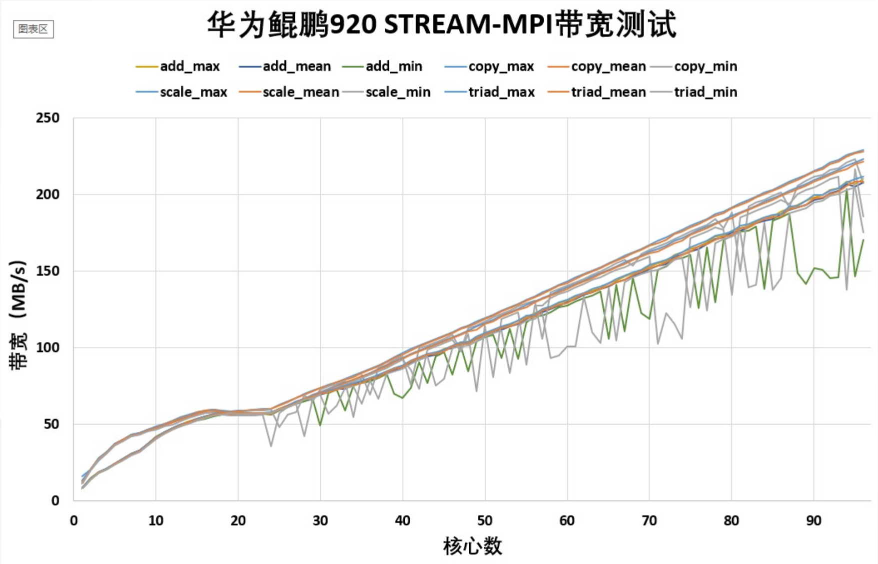

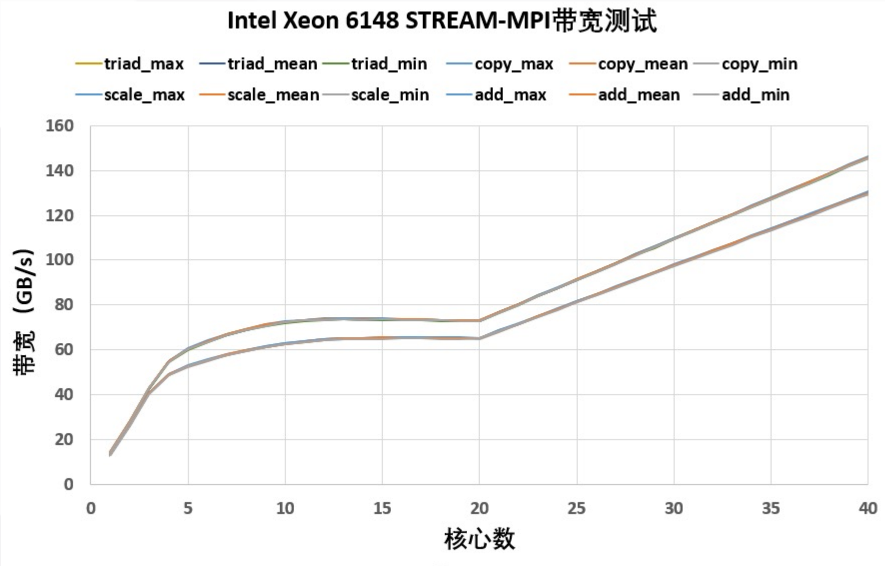

TeaLeaf: Run-to-run performance variation

A further benchmark for Ph.1:

Analyzing the correlations between the bandwidth and performance counters.

Phenomenon：

Both the L1 dCache Refill Rate counter and the Backend Stall counter are linear negative correlated to the DRAM bandwidth.

Conclusion:

1. The multi-core DRAM memory bandwidth is a performance variation source on KP920.
2. The root cause of the performance variation may be related to the mechanism of prefetching and cache eviction.

The relation ship between single-core bandwidth and L1 dCache Refill Rate counter

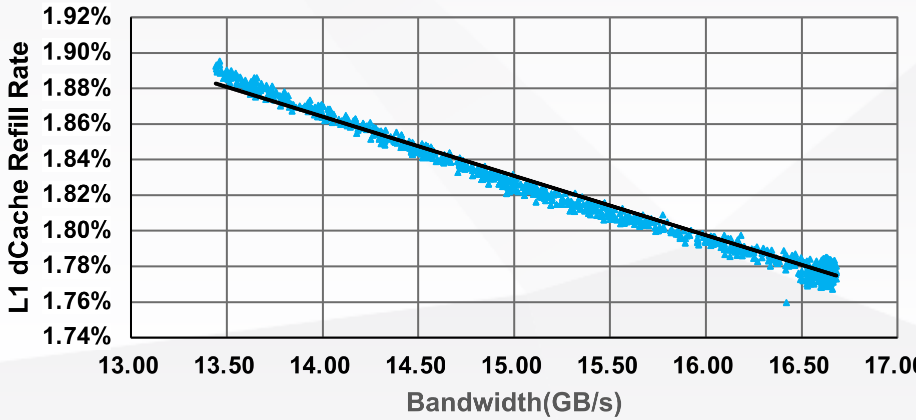

L1 dCache Refill: Attributable instruction memory accesses that cause a refill of at least the Level 1 instruction or unified cache.
L1 dCache Refill Rate: # of L1 dCache Refill / # of L1 dCache Access

The relation ship between single-core bandwidth and Backend Stall counter

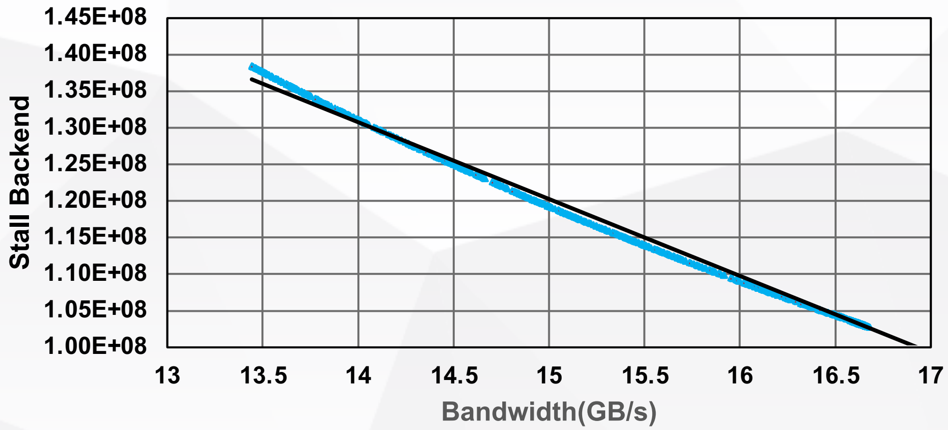

Backend Stall: # of cycles that no operation issued due to backend

SJTU

CloverLeaf: The performance variation in MPI communication

Evaluation:

1. Evaluating P2P bandwidth and latency in MPI communications

Phenomenon：

1. After strong scaling to dual sockets in bm4 case, the performance gap between Xeon6148 and KP920 is enlarged.
2. The large dataset has more performance gap between the two platforms, but less gap between the 1-sock and the 2-sock test.

Further micro benchmarking:

1. Evaluating P2P bandwidth and latency in MPI communications.

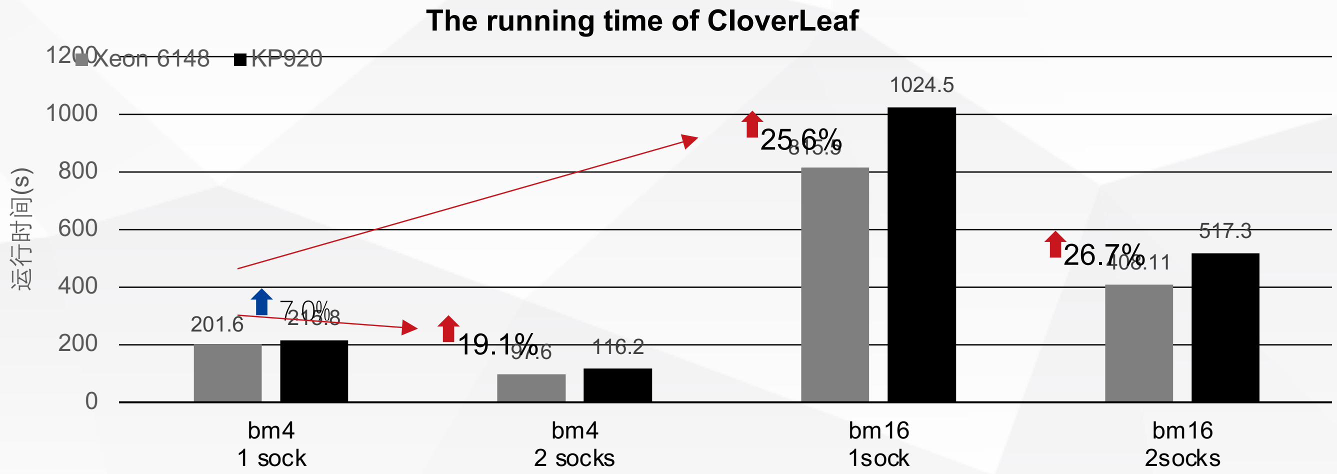

CloverLeaf: Run-to-run performance variation

A further benchmark for Ph.1:

Evaluating P2P bandwidth and latency in MPI communications.

Phenomenon：

1. Compared to Xeon6148, KP920 has higher latency and a steeper increasing trend.

2. Along with the increase in packet sizes, KP920 has a steep and turbulent trend in the MPI bandwidth.

Conclusion:

1. The performance variation in MPI P2P communication leads to inconsistent performance trends during the change of packet sizes in MPI applications.

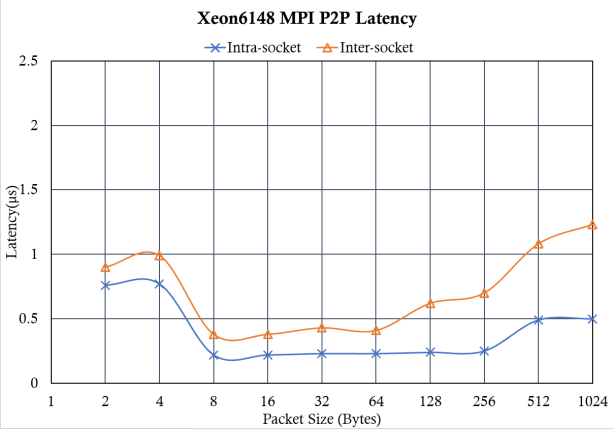

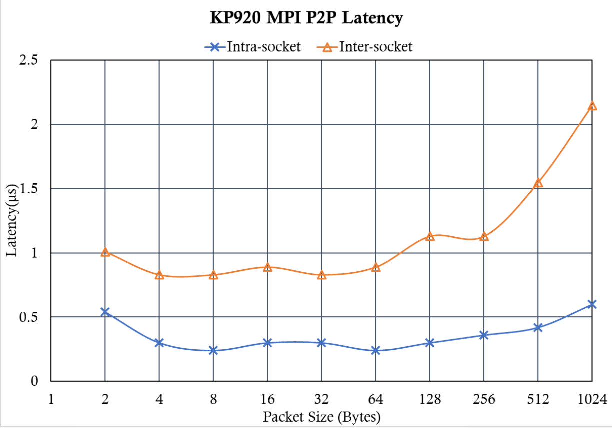

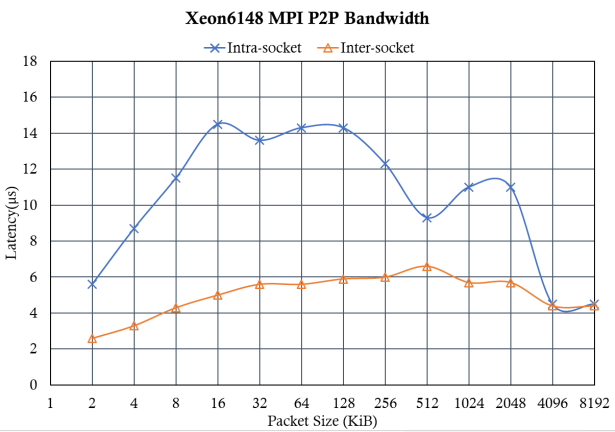

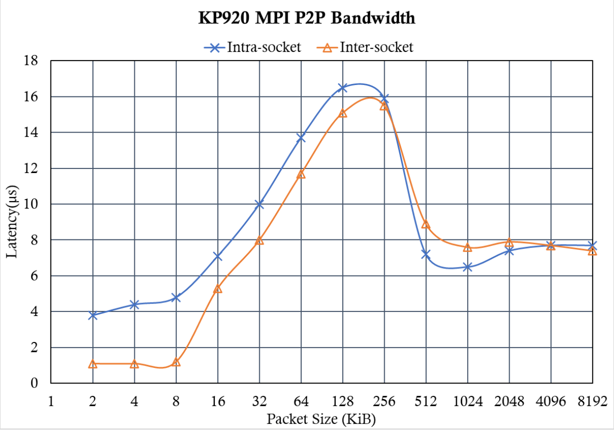

A summary to the performance variation on KP920

Arithmetic instructions

- **Performance variation sources:**

1. The CPI of common arithmetic instructions varies in different precision, and FMA is halved in double precision.

- **Impacts:**

1. The Performance of mix-precision applications will vary in different implementations.
2. Arithmetic with different precision will meet inconsistent performance.

Cache and memory subsystem

- **Performance variation sources:**

1. Variations of the bandwidth and the latency in different compute patterns and different remote access patterns.
2. Run-to-run performance variation in the memory subsystem due to the design in prefetching and cache eviction algorithm.

- **Impacts:**

1. Performance drop when performing inter-CCL and inter-die data sharing.
2. Performance anomaly in memory-bound applications.

MPI communication:

- **Performance variation sources:**

1. Along with the increase in packet sizes, KP920 has a steep and turbulent trend in the MPI bandwidth and latency

- **Impacts:**

1. Inconsistent performance trends during the change of packet sizes in MPI applications.

Outline

SHANGHAI JIAO TONG UNIVERSITY

1 Background

2 The performance variation of KP920

3 Case Study: SJTU KP920 HPC system

4 Conclusion

The SJTU KP920 HPC System

of Nodes: 100 blade nodes.

CPU: KP920（64 cores, 2.6GHz） x 2

Memory: 16GB DDR4 2933MHz x 16

Network: 100Gbps, Mellanox Infiniband EDR

OS: CentOS 7.7 (3.10.0-1062）

Online since May. 2021

Performance evaluation in the use of performance variation

We designed special HPL test cases for triggering performance variation for discovering potential hardware or software failures

A classified statistic of errors in HPC systems of SJTU

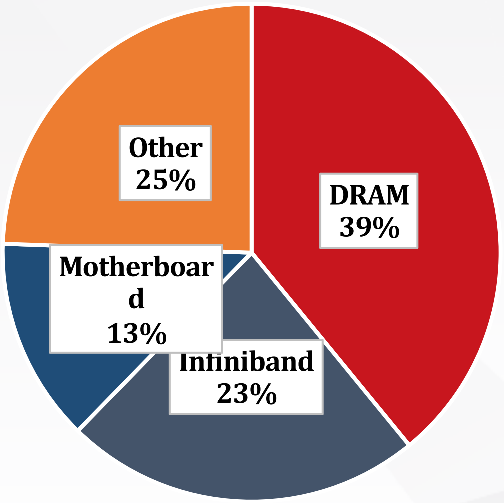

Case	Test parameters	Variation threshold
Single-node HPL-Light	N = 32768, NB = 128, 500 runs	$\frac{1\%\text{quantile}-\text{median}}{\text{median}} > -0.5\%$
32-node HPL-Light	N = 131072, NB = 128, 500 runs per test. Dividing 100 nodes into 3 groups, each group connects to the same IB switch.	$\frac{1\%\text{quantile}-\text{median}}{\text{median}} > -0.25\%$

Performance evaluation in the use of performance variation

In single-node HPL-light test:

Low-performance heavy-tailed results were found.

Conclusion:

Two motherboards malfunction.

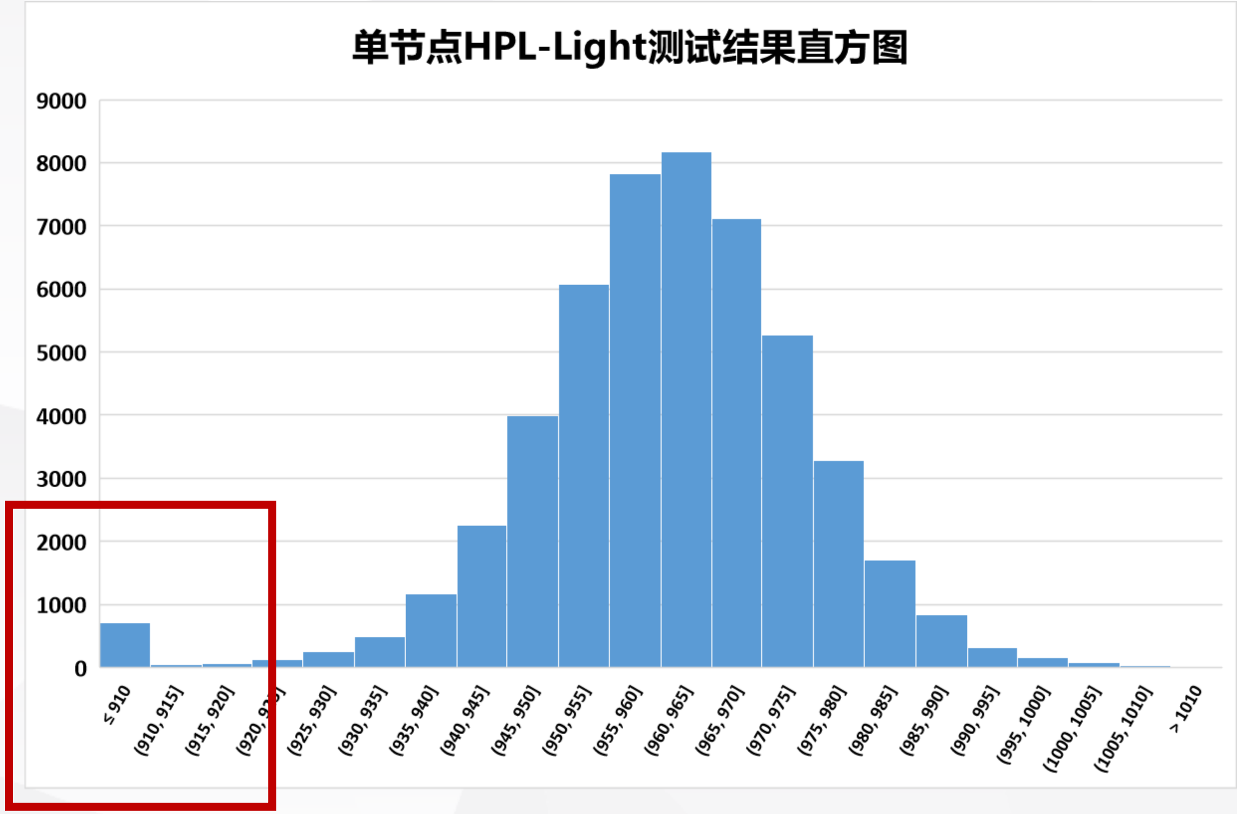

Performance evaluation in the use of performance variation

In 32-node HPL-light test:

Low-performance heavy-tailed results were found in the 32-node HPL-Light test.

Conclusion:

8 out of 100 Infiniband network adapters’ models are not aligned with the others.

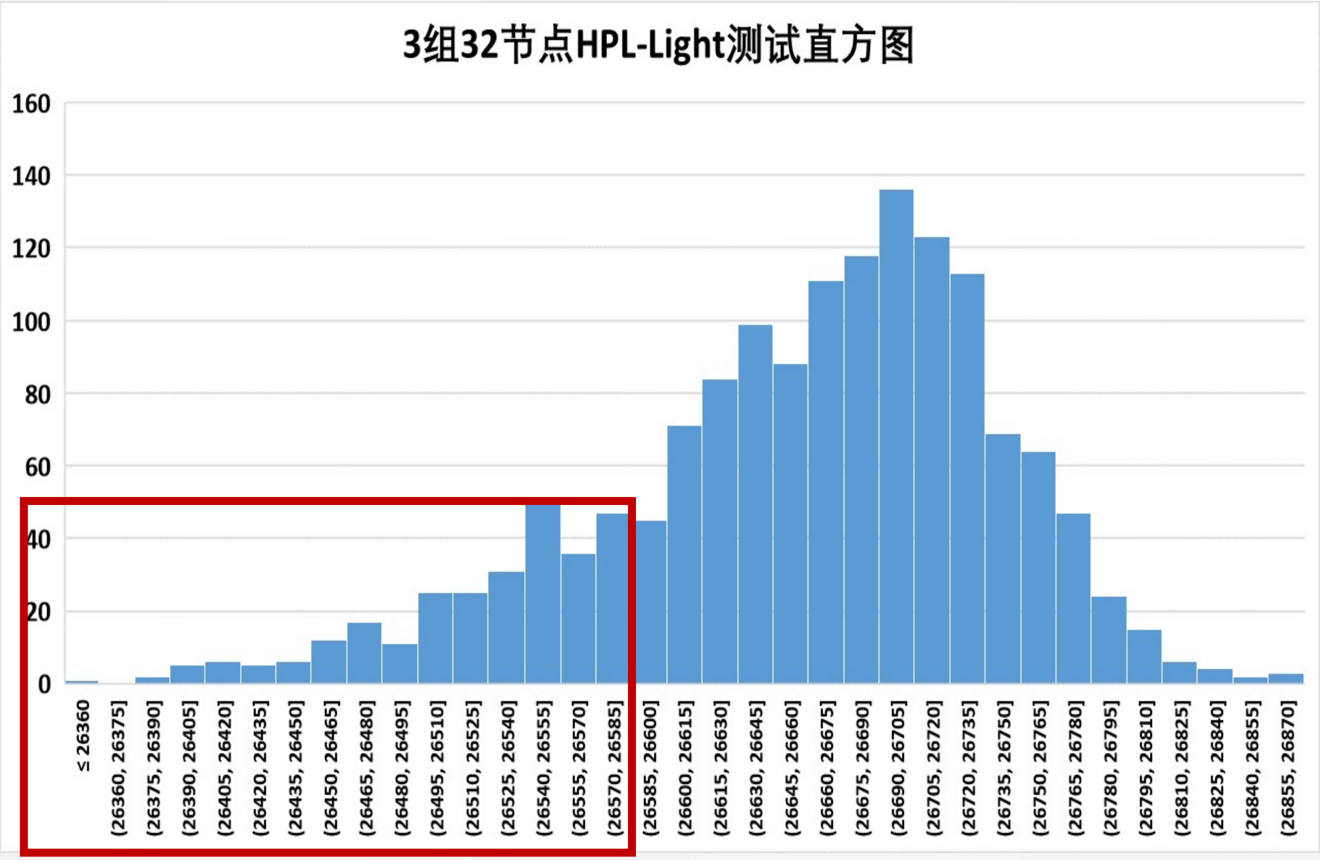

KP920 HPC cluster in production

Malfunction records：

- During building: Replacing 2 motherboards and 8 IB adapters.
- Built till now: Only once hardware replacing。

Utilization ratio: Daily utilization > 75%.

https://account.hpc.sjtu.edu.cn/top

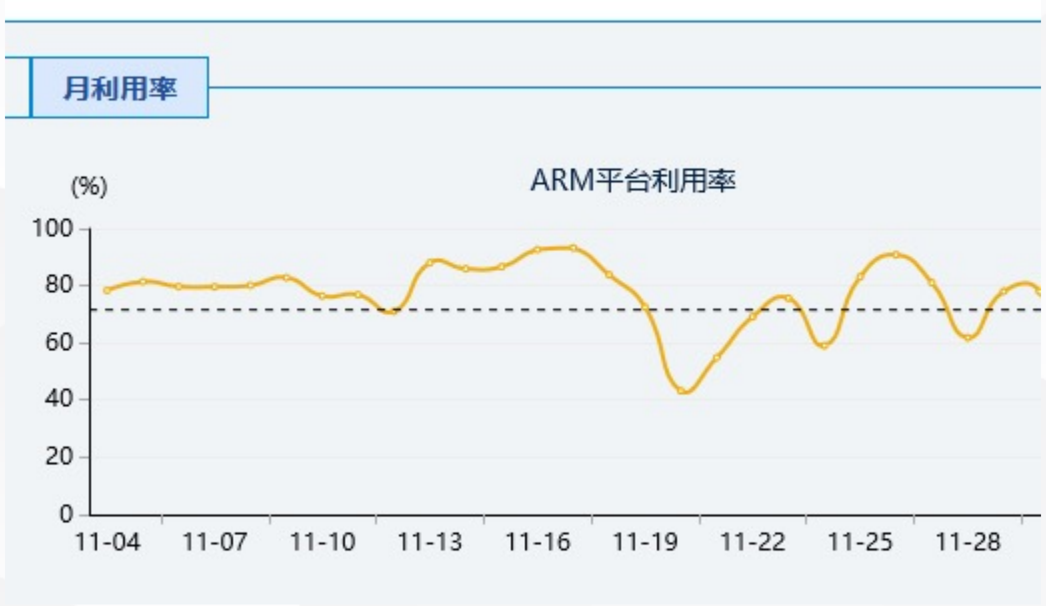

KP920 HPC cluster in production

Part of software in use（https://docs.hpc.sjtu.edu.cn/app/index.html）

GROMACS	A molecular dynamics package.	WRF	Meso-scale weather research forecasting.
VASP	A package for performing ab initio quantum mechanical calculations.	CP2k	Quantum chemistry and solid state physics program package
Amber	A family of force fields for molecular dynamics of biomolecules.	Apache TVM	Compiling framework for machine learning.
LAMMPS	A molecular dynamics program from Sandia National Laboratories.	BCFtools	Tools for bioinformatic vcf/BCF files.
LAMMPS-RBE	A SJTU developed software based on LAMMPS.	Blast-plus	Basic Local Alignment Search Tool.
Manta	Detection of germline mutation and somatic mutation in the tumor/normal coupling sample.	DELLY	Integrated Structural Variant Discovery.
Quantum Espresso	A suite for first-principles electronic-structure calculations and materials modeling.	HISAT2	A fast and sensitive alignment program for mapping next-generation sequencing reads.
		HMMER	Biosequence analysis using profile hidden Markov models.

上海交通大學
SHANGHAI JIAO TONG UNIVERSITY

Outline

The toolset: Top-down measurement and analysis

Application	Description	Usage
HPL	The benchmark that solves an FP64 dense linear system.	Evaluating for double-precision arithmetic performance.
STREAM	A memory streaming benchmark.	Evaluating for DDR memory performance.
OSU Benchmark	A high-speed interconnect benchmark suite.	Evaluating MPI communication performance.
SNAP	A mini-app benchmark that simulates the solving of linear Boltzmann transport equation (TE)。	Evaluating the performance of memory bandwidth and cache latency.
TeaLeaf	A mini-app benchmark that solves the linear heat conduction equation.	Evaluating the performance of regular memory access in stencil compute.
CloverLeaf	A mini-app benchmark that solves the compressible Euler equations in 2D	Evaluating the performance of communication between neighbor processes, streaming bandwidth, and FP64 arithmetic.
PAPI、LIKWID	The tool for µarch benchmarking, monitoring, and data collection.	Collecting performance data and locating the sources of processor performance variations.

IEEE Cluster 2024

Shanghai, China

September 2024

General co-chairs

Dr. Satoshi Matsuoka
- Director, RIKEN-CCS
- SC13 Program Chair

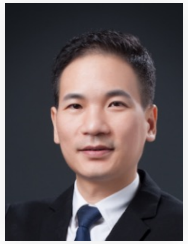

Dr. James Lin
- Vice Director, HPC Center at SJTU
- Cluster SC member

TPC co-chairs

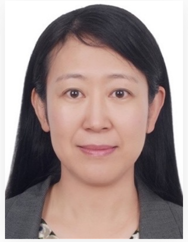

Dr. Yutong Lu
- Director, National Supercomputer Center in Guangzhou
- ISC19 Program Chair

Dr. Wu-chun Feng
- Professor, Virginia Tech
- SC13 TPC Chair