

Summit System Overview

Tom Papatheodore

Oak Ridge Leadership Computing Facility (OLCF)
Summit New User Training

June 3, 2020

ORNL is managed by UT-Battelle, LLC for the US Department of Energy



U.S. DEPARTMENT OF
ENERGY

OLCF Summit Overview

The system includes


- 4,608 nodes
- Dual-port Mellanox EDR InfiniBand network
- 250 PB IBM file system transferring data at 2.5 TB/s

System Performance

- Peak of 200 Petaflops (FP₆₄) for modeling & simulation
- Peak of 3.3 ExaOps (FP₁₆) for data analytics and artificial intelligence

Each node has

- 2 IBM POWER9 processors
- 6 NVIDIA Tesla V100 GPUs
- 608 GB of fast memory (96 GB HBM2 + 512 GB DDR4)
- 1.6 TB of NV memory




HPCC List for November 2019

Rank	System	Cores	Rmax (TFlop/s)	HPCC (TFlop/s)
1	Summit - IBM Power System AC922, IBM POWER9 22C 3.07GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR InfiniBand, IBM DOE/SC/Oak Ridge National Laboratory, United States	2,414,592	148,600.0	2925.75




Summit Node Schematic



IBM Power9 Processor


- Summit's P9s: 22 cores (4 hwthreads/core)
- PCI-Express 4.0
 - Twice as fast as PCIe 3.0
- NVLink 2.0
 - Coherent, high-bandwidth links to GPUs
- 14nm FinFET SOI technology
 - 8 billion transistors
- Cache
 - L1I: 32 KiB (per core, 8-way set associative)
 - L1D: 32 KiB (per core, 8-way)
 - L2: 512 KiB (per pair of cores)
 - L3: 120 MiB eDRAM, 20-way (shared by all cores)



Summit Contains 27,648 NVIDIA Tesla v100s

Each Tesla v100 GPU has:

- 150+150 GB/s total BW (NVLink v2.0)
- 5,120 CUDA cores (64 on each of 80 SMs)
- 640 Tensor cores (8 on each of 80 SMs)
- 20MB Registers | 16MB Cache | 16GB HBM2 @ 900 GB/s
- 7.5 DP TFLOPS | 15 SP TFLOPS | 120 FP₁₆ TOPS
- Tensor cores do mixed precision multiply-add of 4x4 matrices



$$\mathbf{D} = \begin{pmatrix} \begin{matrix} A_{0,0} & A_{0,1} & A_{0,2} & A_{0,3} \\ A_{1,0} & A_{1,1} & A_{1,2} & A_{1,3} \\ A_{2,0} & A_{2,1} & A_{2,2} & A_{2,3} \\ A_{3,0} & A_{3,1} & A_{3,2} & A_{3,3} \end{matrix} & \begin{matrix} B_{0,0} & B_{0,1} & B_{0,2} & B_{0,3} \\ B_{1,0} & B_{1,1} & B_{1,2} & B_{1,3} \\ B_{2,0} & B_{2,1} & B_{2,2} & B_{2,3} \\ B_{3,0} & B_{3,1} & B_{3,2} & B_{3,3} \end{matrix} \\ \text{FP16 or FP32} & \text{FP16} \end{pmatrix} + \begin{pmatrix} \begin{matrix} C_{0,0} & C_{0,1} & C_{0,2} & C_{0,3} \\ C_{1,0} & C_{1,1} & C_{1,2} & C_{1,3} \\ C_{2,0} & C_{2,1} & C_{2,2} & C_{2,3} \\ C_{3,0} & C_{3,1} & C_{3,2} & C_{3,3} \end{matrix} \\ \text{FP16 or FP32} \end{pmatrix}$$

$$\mathbf{D} = \mathbf{AB} + \mathbf{C}$$

Type	Size	Range	$u = 2^{-t}$
half	16 bits	$10^{\pm 5}$	$2^{-11} \approx 4.9 \times 10^{-4}$
single	32 bits	$10^{\pm 38}$	$2^{-24} \approx 6.0 \times 10^{-8}$
double	64 bits	$10^{\pm 308}$	$2^{-53} \approx 1.1 \times 10^{-16}$
quadruple	128 bits	$10^{\pm 4932}$	$2^{-113} \approx 9.6 \times 10^{-35}$


https://docs.olcf.ornl.gov/systems/summit_user_guide.html#tensor-cores

Summit Node

(2) IBM Power9 + (6) NVIDIA Volta V100




Summit Board (1 node) showing the Water Cooling



Summit Specs

summit



Feature	Summit
Peak FLOPS	200 PF
Max possible Power	13 MW
Number of Nodes	4,608
Node performance	42 TF
Memory per Node	512 GB DDR4 + 96 GB HBM2
NV memory per Node	1.6 TB
Total System Memory	2.8 PB + 7.4 PB NVM
System Interconnect	Dual Port EDR-IB (25 GB/s)
Interconnect Topology	Non-blocking Fat Tree
Bi-Section Bandwidth	115.2 TB/s
Processors on node	2 IBM POWER9™ 6 NVIDIA Volta™
File System	250 PB, 2.5 TB/s, GPFS™

Questions?

Summit here



Frontier here

