



Greenest Supercomputer in the World

TSUBAME-KFC Overview

TSUBAME-KFC system is a prototype for future generation supercomputers such as TSUBAME 3.0 and beyond. Main purpose is to optimize the power efficiency, both compute nodes and cooling system.

- 40+ compute nodes connected via FDR InfiniBand network
- two Intel Xeon E5-2620 V2 Processor (Ivy Bridge EP)
- four NVIDIA Tesla K20X GPUs
- 64GB DDR3-1600 Memory, SATA SSDs (120+480+480GB)



Certificates of #1 in Green500 (Nov. 2013, Jun. 2014) and GreenGraph500 (Nov. 2013). The TSUBAME-KFC system is ranked #1 in two energy efficiency rankings.



The container-based server room and cooling tower are located beside our computing center building.




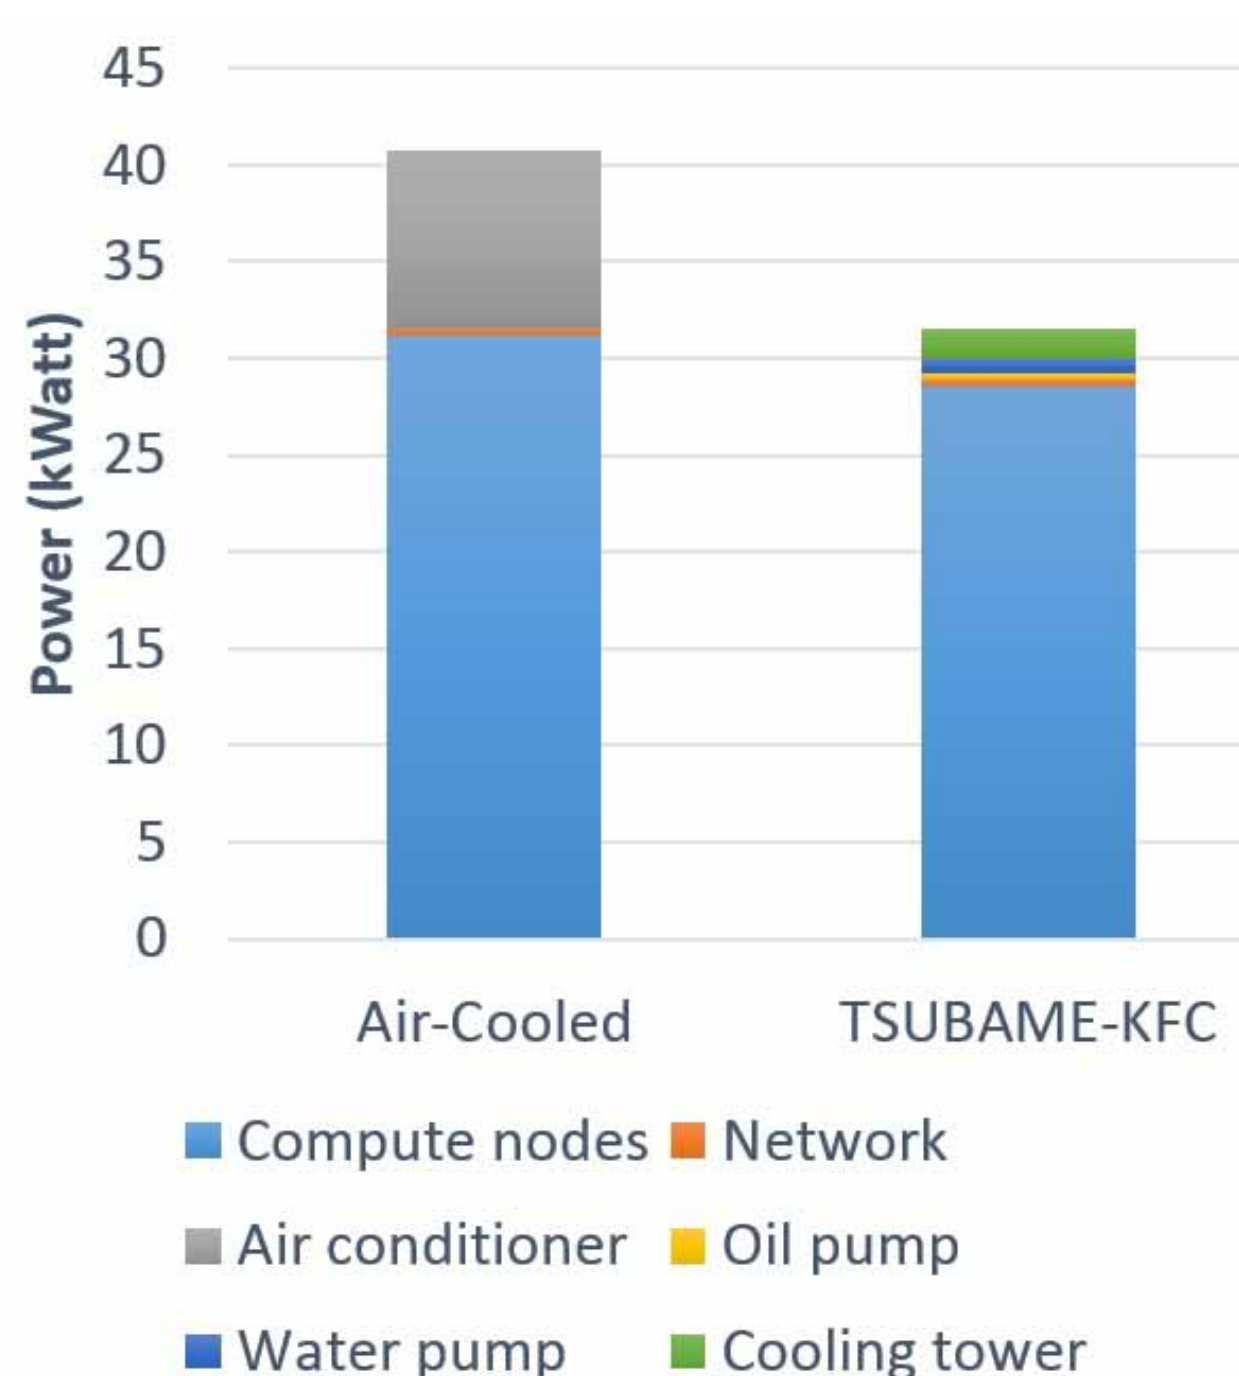
Oil-submerged compute nodes



A compute node after oil-submerged

Power Usage Effectiveness

Power Usage Effectiveness (PUE) is one of the major metric to show the power efficiency of cooling system. The PUE of TSUBAME 2.0/2.5 is about 1.29, which means additional 29% power is used for cooling system. This value is better than those of standard datacenters. PUE of TSUBAME-KFC is less than 1.1, thanks to oil-submersion cooling technology. TSUBAME-KFC uses evaporation cooling tower instead of chiller, as a result the necessary power is only for pump systems (for oil and water), and fans in the cooling tower.



Many power meters are installed to measure the power consumption of each node, network switch, pump, and cooling tower.



Hardware Maintenance

The compute nodes are submerged in the oil. It increases the cost of hardware maintenance in case of trouble. The table on the right side shows the breakdown of the time spent in each procedure when installing additional SSD drivers to each node. This operation requires at least two engineers to lift up the node safely.

Due to the strong cooling, we never observe overheating errors both CPU and GPU. The temperature of GPUs are basically kept under 50 degrees C.

Procedure	Time
Removing external cables	50s
Pulling up the node	2m10s
Opening chassis cover	1m10s
Removing GPUs	1m20s
Installing SATA cables	1m50s
Restoring GPUs	2m20s
Closing chassis cover	2m40s
Inserting SSDs into drive bay	1m10s
Submerging the nodes	2m30s
Installing external cables	40s
Total	16m40s

[1] Toshio Endo, Akira Nukada and Satoshi Matsuoka. "TSUBAME-KFC: a Modern Liquid Submersion Cooling Prototype towards Exascale Becoming the Greenest Supercomputer in the World", In Proc. of the 20th IEEE International Conference on Parallel and Distributed Systems (ICPADS'14), Dec. 2014, to appear.

Collaboration with NEC, Green Revolution Cooling, NVIDIA, SuperMicro. Supported by MEXT.