

Report on SC13 joint Green500/EE HPC WG Power Measurement Methodology BoF

TITLE: The Green500 List and Its Evolution

BoF Moderator:

- Wu Feng, Virginia Tech

Presenters:

- Erich Strohmaier, Lawrence Berkeley National Laboratory (LBNL)
- Robin Goldstone, Lawrence Livermore National Laboratory (LLNL)
- Thomas Schulthess, Swiss Supercomputing Center (CSCS)
- Satoshi Matsuoka, Tokyo Institute of Technology (TITech)

This BoF encompassed the following:

- Presentation and discussion of (1) the evolving metrics, methodologies, and workloads for energy-efficient HPC, (2) the trends across the Green500, including the trajectory towards exascale, and (3) highlights from the latest Green500 List.
- Reports on a collaborative effort between the Green500, Top500, and the EE HPC WG to improve power-measurement methodology while running a workload, such as High Performance LINPACK (HPL). In particular, Robin Goldstone of LLNL and Thomas Schulthess of CSCS presented their respective results from using the improved methodology at their sites.
- Presentation of the Green500 awards, including the greenest supercomputer, namely Tsubame-KFC at 4.5 gigaflops/watt, and the greenest petaflop supercomputer, namely Piz Diant at the aforementioned CSCS.
- An overview of Tsubame-KFC, the greenest supercomputer on the Green500, by Satoshi Matsuoka of the Tokyo Institute of Technology.

The BoF attendance and participation was strong, with well over 100 people in attendance. It was advertised through the Green500 website and the EE HPC WG website, at the SC13 Workshop on “Building Energy Efficient HPC,” and at the EE HPC WG Booth on the SC13 Exhibition Floor.

Presentations from the BoF can be found at the following web site:

<http://eehpcwg.lbl.gov/documents/sc13-technical-program/birds-of-feather>

Additional discussion and feedback from BoF participants:

- A few clarifying questions were asked about the different power-measurement methodologies: Level 1 (“good”), Level 2 (“better”), and Level 3 (“best”).
- Questions were asked about the different energy-efficiency numbers between Level 1 and Level 2 for LLNL and Level 1 and Level 3 for CSCS. The differences were largely attributed to the network subsystem, which in the past consumed a smaller fraction of the overall power consumption. However, for the system at CSCS, the network consumed a much larger portion of the overall power consumption.

- Debate about what measurements the Green500 should be accepting, in particular, Level 1 vs. Level 2 or Level 3.
 - History: Level 1 was set-up to encourage measurement with fairly high accuracy (as long as the network subsystem did not consume a large fraction of the overall power consumption). However, even with this lowered bar of measurement entry, only slightly less than 300 machines reported their power; the remaining 200+ machines had their power numbers derived, resulting in a less competitive efficiency rating (which was intended to incentivize a measured submission in subsequent lists, which it did to some extent).
 - This Year:
 - Green500 accepts higher-precision measurements, denoted as Level 2 and 3, as a supplement to Level 0 (derived) and Level 1.
 - Only *six* Level 2/Level 3 submissions have been made to the Green500 over the course of the past year.
 - Evolution:
 - Continue to accept and encourage Level 1, Level 2, and Level 3 submissions on the Green500 and potentially evolve it so that the Level 2 measurement is the default measuring stick of comparison.

GOAL

- *“Higher quality measurements... provide a much better picture of the real-world costs... as well as a more in-depth picture of how the system handles a Linpack run.”*
Green500 press release