

## Advanced Computing, Mathematics and Data Division Research Highlights

April 2015

### Energy Star

#### *Novel models of HPC systems depict the interplay between energy efficiency and resilience*

**Results:** To achieve high resiliency and efficiency on future exascale computing systems (capable of a billion, billion calculations per second), scientists from Pacific Northwest National Laboratory; University of California, Riverside; and Marquette University examined some advanced high-performance computing systems and determined that undervolting, or dynamic voltage scaling to reduce power consumption, that leverages existing mainstream resilience techniques at scale improved system failure rates. The model also demonstrated up to 12 percent energy savings over baseline runs (with eight HPC benchmarks) and up to nine percent savings against state-of-the-art dynamic voltage and frequency scaling (DVFS) solutions currently used to lower the operating frequency (supply voltage changes according to frequency) of hardware. Notably, the results are based on a conservative assumption of the total energy savings because the model applies the peak range of the failure rates.

The paper documenting this first-of-its-kind work, “Investigating the Interplay between Energy Efficiency and Resilience in High Performance Computing,” was accepted by the 29th IEEE International Parallel & Distributed Processing Symposium (IPDPS), a premier international forum that showcases the latest research findings in all aspects of parallel and distributed computing, and will be presented during the IPDPS 2015 conference’s technical program in a session devoted to resilience on May 27, 2015.

**Why it Matters:** While microprocessors can run faster, the heat they generate, especially as they move voluminous amounts of data, makes it prohibitive. To break through to true exascale-level computing ( $10^{18}$  flops) operations, the heat output from computing systems must be curtailed.

Moreover, to do so within a 20-megawatt power wall (limit) means improving on current supercomputing systems that, merely to operate, can consume enough energy to power a small city or thousands of homes.

“Power sources are not unlimited nor are they free, and in a computing system, primary reasons for failures include radiation from the cosmic rays, packaging materials, and temperature fluctuation. These are crucial challenges affecting the push to extreme-scale systems,” said [Shuaiwen Leon Song](#), a research scientist with PNNL’s High Performance Computing group and a co-author of the paper describing the research. “Our undervolting method does not modify existing hardware or require pre-production machines and has shown positive results toward achieving a cost-efficient energy-savings implementation for the HPC field.”

**Methods:** For their work, the researchers had to determine if the trade-off between power savings through

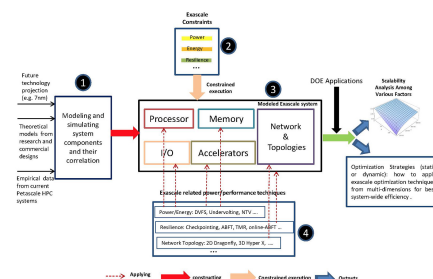


Figure depicting the general constraints at scale for building exascale architectures. The undervolting component (shown in 4) is part of ongoing work that PNNL’s HPC group and their collaborators are conducting in the effort to build highly resilient and energy-efficient HPC systems. [Enlarge Image.](#)

undervolting and performance overhead for coping with higher failure rates could reduce overall energy consumption. In the process, they would clarify if future exascale systems should trend toward using low-voltage embedded architectures for energy-efficiency or rely primarily on advanced software-level techniques to achieve high system resiliency and efficiency. Unfortunately, undervolting on its own often results in increased hard (e.g., system abort from power outage) and soft (e.g., memory bit flips) errors, diminishing its viability as a power-saving technique for HPC systems. So, targeting general faults on common HPC production machines at scale, they examined performance and energy efficiency of several HPC runs with undervolting and different mainstream resilience techniques on power-aware clusters. In addition, they examined normalized performance and energy efficiency of a current DVFS technique, Adagio, a runtime system that provides energy savings in HPC applications, both as a standalone option and combined with undervolting. Using general power and energy models afforded a closer look at the factors that likely affect overall energy savings through undervolting, as well as the interplay between power reduction through undervolting and application performance loss that stems from the required fault detection and recovery at elevated failure rates caused by undervolting.

The conclusions from their extensive experimental results indicate that software-level techniques can play an important role for future exascale systems to address the complex entangled effects caused by energy and resilience at scale.

**What's Next?** Exploring undervolting is another aspect of ongoing research within PNNL's HPC group seeking solutions to challenges involving energy and performance in future extreme-scale computing systems.

**Acknowledgments:** The research was supported by the U.S. Department of Energy's [Office of Advanced Scientific Computing Research](#) (PNNL's Beyond the Standard Model project), as well as the National Science Foundation through grants #CCF-1305622 and #ACI-1305624.

**Reference:**

Tan L, SL Song, P Wu, Z Chen, R Ge, and DJ Kerbyson. 2015. "Investigating the Interplay between Energy Efficiency and Resilience in High Performance Computing." In *29th IEEE International Parallel & Distributed Processing Symposium (IPDPS 2015)*. May 25-29, 2015, Hyderabad, India.

**Related:**

- [PNNL HPC Staff Take on Energy Efficiency, Resilience in New Paper](#)
-