

Dynamic Power Management for Megawatt-Sized Supercomputer Centers

Tapasya Patki
LLNL & The University of Arizona

 Lawrence Livermore
National Laboratory

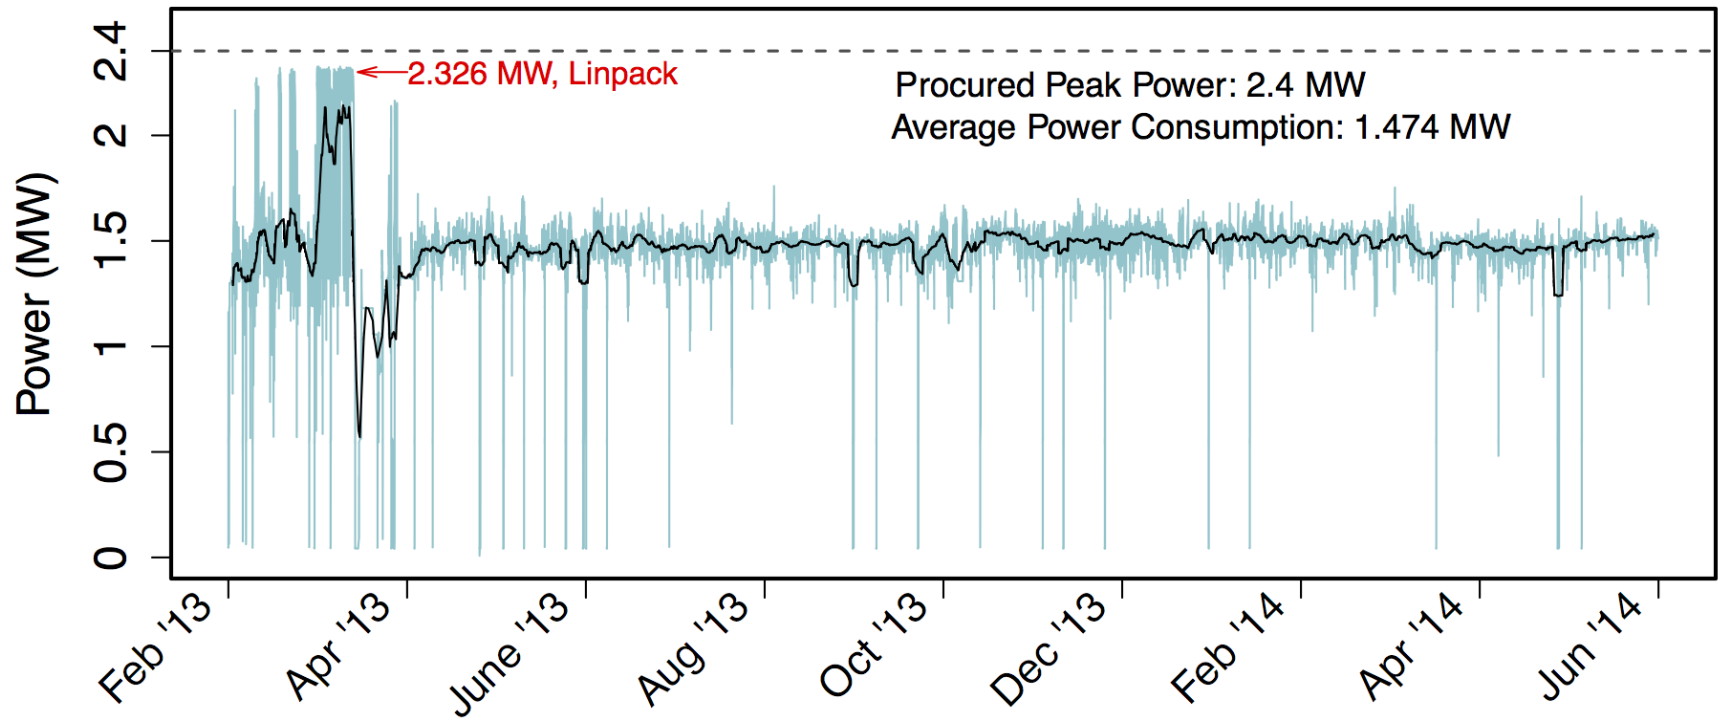
LLNL-PRES-663772

This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344. Lawrence Livermore National Security, LLC



Unused Power

**Total Power Consumption of the BG/Q Vulcan Supercomputer
Feb 2013 to Jun 2014 (3 minute time samples)**



Problem

20

Megawatts

0

↕ 5% Desirable range

Problem is **power utilization**,
not power procurement

Don't:

- “Save” Power
(or, Energy)

Do:

- Use power to do
more science
- Improve application
performance

100%

Power Utilization

0%

Solution: HW Overprovisioning

20

Megawatts

0

↕ 5% Desirable range


- Rethink utilization in terms of power, not nodes
- Not all nodes can run at peak power simultaneously
- Reconfigure dynamically based on workload characteristics to improve performance

100%

Power Utilization

0%

Improving Performance



CFD solver kernel

Traditional configuration

Poor performance,
doesn't use allocated
power

Best configuration

2x improvement in
performance, uses all of
the allocated power

Improving Performance

Choose the **right configuration** under a power bound based on **application characteristics**

- Scalability:
 - Fewer nodes, higher power per node
 - More nodes, lower power per node
- Memory intensity:
 - Vary cores per node

Dynamic Power Management

Site

- Demand Response, Renewables

Cluster

- Overprovisioning, Job scheduling

Job/Application

- Adaptive runtimes, Power balancing


Node

- Measurement & control, Power capping

Inherited Power Bounds




Static global power bound




Renewable energy sources
may lead to
varying global power bounds...

Static global power bound



Renewable energy sources
may lead to
varying global power bounds



Solution:
Dynamically Reconfigurable,
Hardware-Overprovisioned Systems

Thanks

- David Lowenthal, University of Arizona
- Ghaleb Abdulla, LLNL
- Barry Rountree, LLNL

