Measuring Power Consumption on IBM Blue Gene/Q

ILLINOIS INSTITUTE OF TECHNOLOGY

Introduction

- As the field of supercomputing continues to push towards the exascale era, power consumption is becoming an increasingly vital area of research.
- To achieve exascale current supercomputers will need to scale their performance by 160X while increasing their power consumption by just 12X.
- Fortunately, hardware manufacturers are starting to deploy various sensors on HPC systems to collect power related data as well as providing relatively easy access to the data they collect.
- Two ares of collection:
  * "Environmental" database
  * Profiling code built on vendor supplied APIs

Blue Gene/Q Architecture and Environmental Data Collection

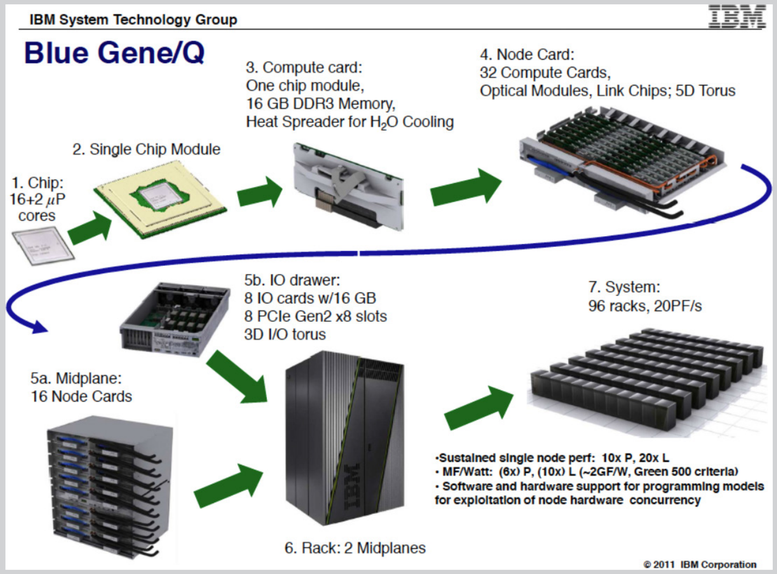

- Environmental database:
  * IBM DB2 relational database.
  * Periodically samples and gathers environmental data from various sensors.
  * Relatively long polling intervals (about 4 minutes on average but can be 60-1800 seconds).
  * While a shorter polling interval would be ideal, the volume of data and stress would exceed the server's processing capacity.

Analysis of Environmental Power and Temp

- One month sample from September 2012
  * Part of the system was undergoing stability and short testing, rest of the system was in maintenance.

Number of Racks	Number of Jobs
1	1,308
2	539
4	318
8	328
16	90
24	1
32	6
48	4
Total	**2,594**

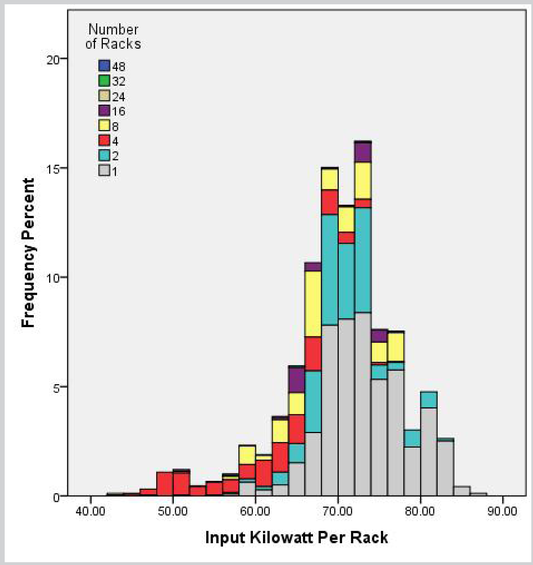

- Most jobs fall into the 60-80 kW bins and all jobs are between 40 and 90 kW.

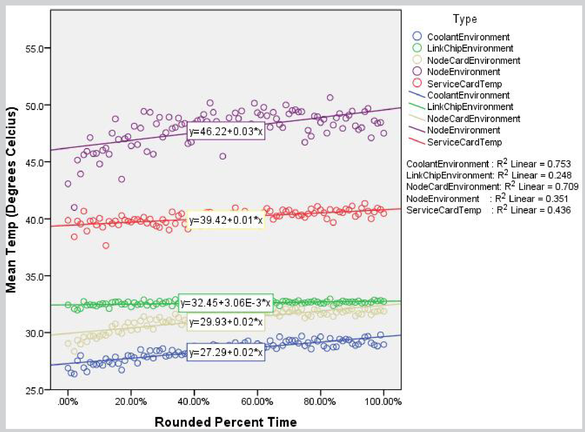

- System gets "hotter" as jobs run.
- Most sensors indicate 2 to 3 degree increase of individual components.

Analysis of EMON Data

- Our profiling code utilizes the EMON API.
- Timer periodically invokes an "EMON system call" at an interval lower than the FPGA interval (i.e., 500ms), records the instantaneous power usage (Watts) across all available domains along with a timestamp.

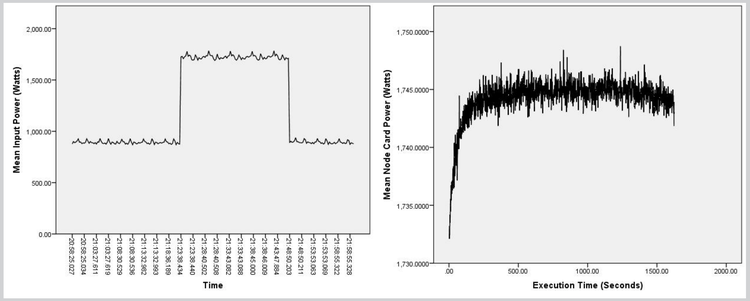

- Almost identical power usage numbers from profiling code and environmental database.
- As there are significantly more data points in the data obtained from our profiling code, it is the more "precise" of the two.
- Higher level of detail afforded by our profiling code means we can draw much more precise conclusions on energy consumption especially as a function of time.

Team Members

- Sean Wallace (Ph.D. Student)
- Zhiling Lan (Faculty)

Collaborators

- Venkatram Vishwanath
- Mike Papka (Argonne)
- Susan Coghlan (Argonne)
- Cheetah Goletz (Argonne)
- Paul Coteus (IBM)
- Christopher M. Marroquin (IBM)

Acknowledgements

This research has been funded in part and used resources of the Argonne Leadership Computing Facility at Argonne National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under contract DEAC02-06CH11357.