

Power and Energy

Scalable Tools Workshop
Working Group Report

Rashawn Knapp, William Jalby, Thiago Teixeira, Allen Malony


Working Group Charter

- Provide guidance to application users/ developers about the scaling limits of an application in its current form with respect to power, energy and the constraints of a machine
- Discuss what we know performance tools currently do in this space

Dummies Guide to Power and Energy

Energy


- Quantitative property that must be transferred to an object in order to perform work on, or to heat, the object
- SI unit of energy is the *joule*, the energy transferred to an object by the work of moving it a distance of 1 meter against a force of 1 newton



Joule Apparatus: measures mechanical equivalent of heat by descending a weigh into water causing a paddle to rotate.

Power:

- The rate of doing work.
- SI unit of power = 1 joule/second (Watt)
- 1 *hp* may be easier to understand: the power required to lift 75 Kg by one meter in 1 second



By wwn work, original version in German by Sgbeer - File:Pferdestaerke.svg, CC BY-SA 3.0, <https://commons.wikimedia.org/w/index.php?curid=35217113>

Problem Definition

- Assume you have a limited power budget
 - On a system platform
- Assume you are running a single application
- How large can an application scale before it exceeds the power budget?
- What do you do then to enable further scaling?

What is needed?

- Understand power usage over
 - Lifetime of execution
 - Program structure
 - Routine level
 - Computing versus communication
 - Hardware components and variety
 - Different code version (e.g., uses of vector instructions and length)
- Need tools to:
 - Measure power consumption
 - Model power usage
 - Control power consumption
 - Manage power usage
- Need program transformation to produce variants
 - Variants have different power requirements
 - Variants may affect compute time (performance)

Tools that We Know Of

- RAPL (Me,C)
- SEP / EMON (Me)
- READEX (Me,C,Ma)
- POW (Ma) (uses RAPL)
- Linux Governor (Ma)
- ...

Me: measurement

C: control

Ma: management

Just to Be Clear

- Can not exceed instantaneous power limit
- It is ok to run longer to prevent this
- What is done when the limits are met?

Recommendation

- More tools are needed and deployed
- Need more talks about energy
- Performance does not mean only time
- ...