





On Chip Controller(OCC) Overview

Todd Rosedahl, Chief Engineer
IBM/POWER Firmware Development




Agenda




- Introduction/Motivation
- System Stack/Ecosystem
- Video
- OCC Overview
- Hardware Block Diagram
- Firmware Details
- Sample Results – Graphs
- Future Enhancements
- References

Motivation



- US datacenter energy consumption – 91Billion KWH
- 34 500MW power plants
- State of Minnesota – 68Billion KWH
- Power Usage Effectiveness

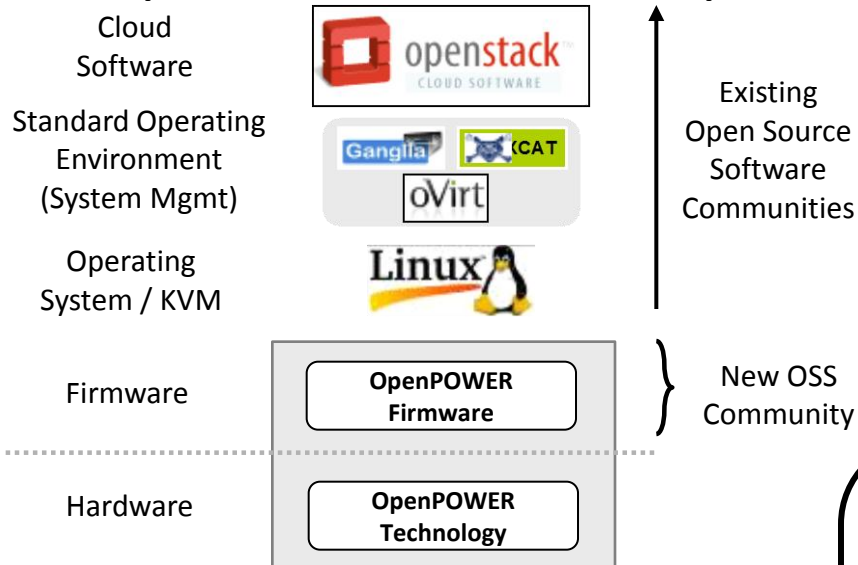
Comparison of Projected Electricity Use [All Scenarios 2007-2011]



Ecosystem Enablement




Power Open Source Software Stack Components



July 2014 Power8 open source firmware stack contributed thru GitHub

Toolkits and resources for porting and optimizing, growing repository on website

Firmware Technical Training – OpenPower Pavilion
Thursday 9-11am Architectural Overview
Thursday 12-2pm Hands-on Training




Your logo
here



Join the conversation at [#OpenPOWERSummit](https://twitter.com/OpenPOWERSummit)

OCC Overview

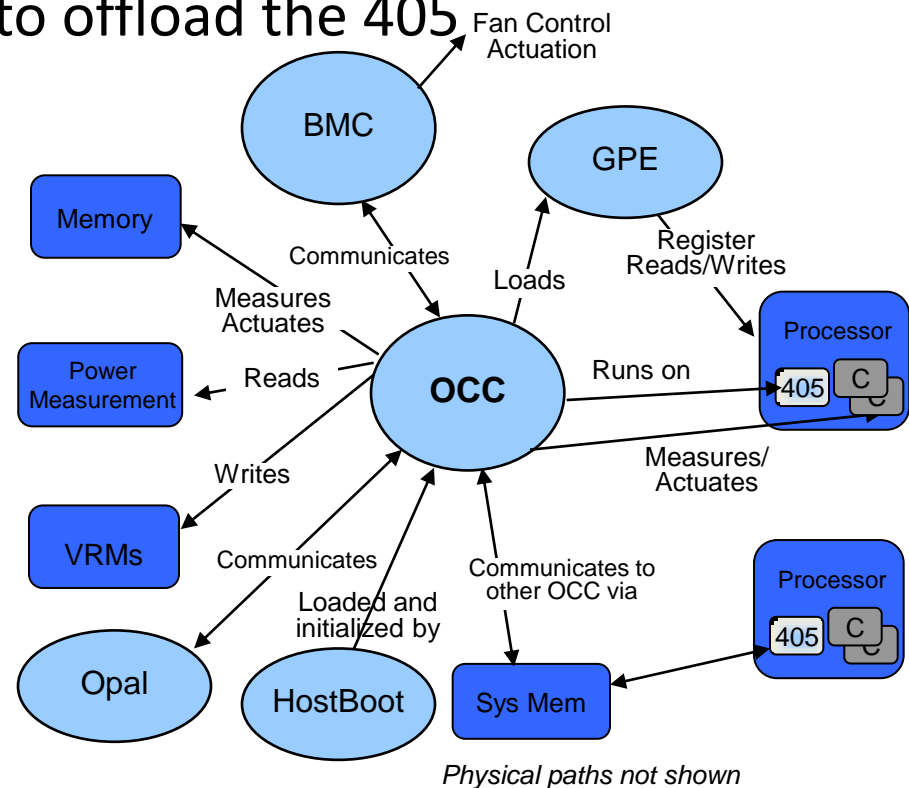


What is OCC?

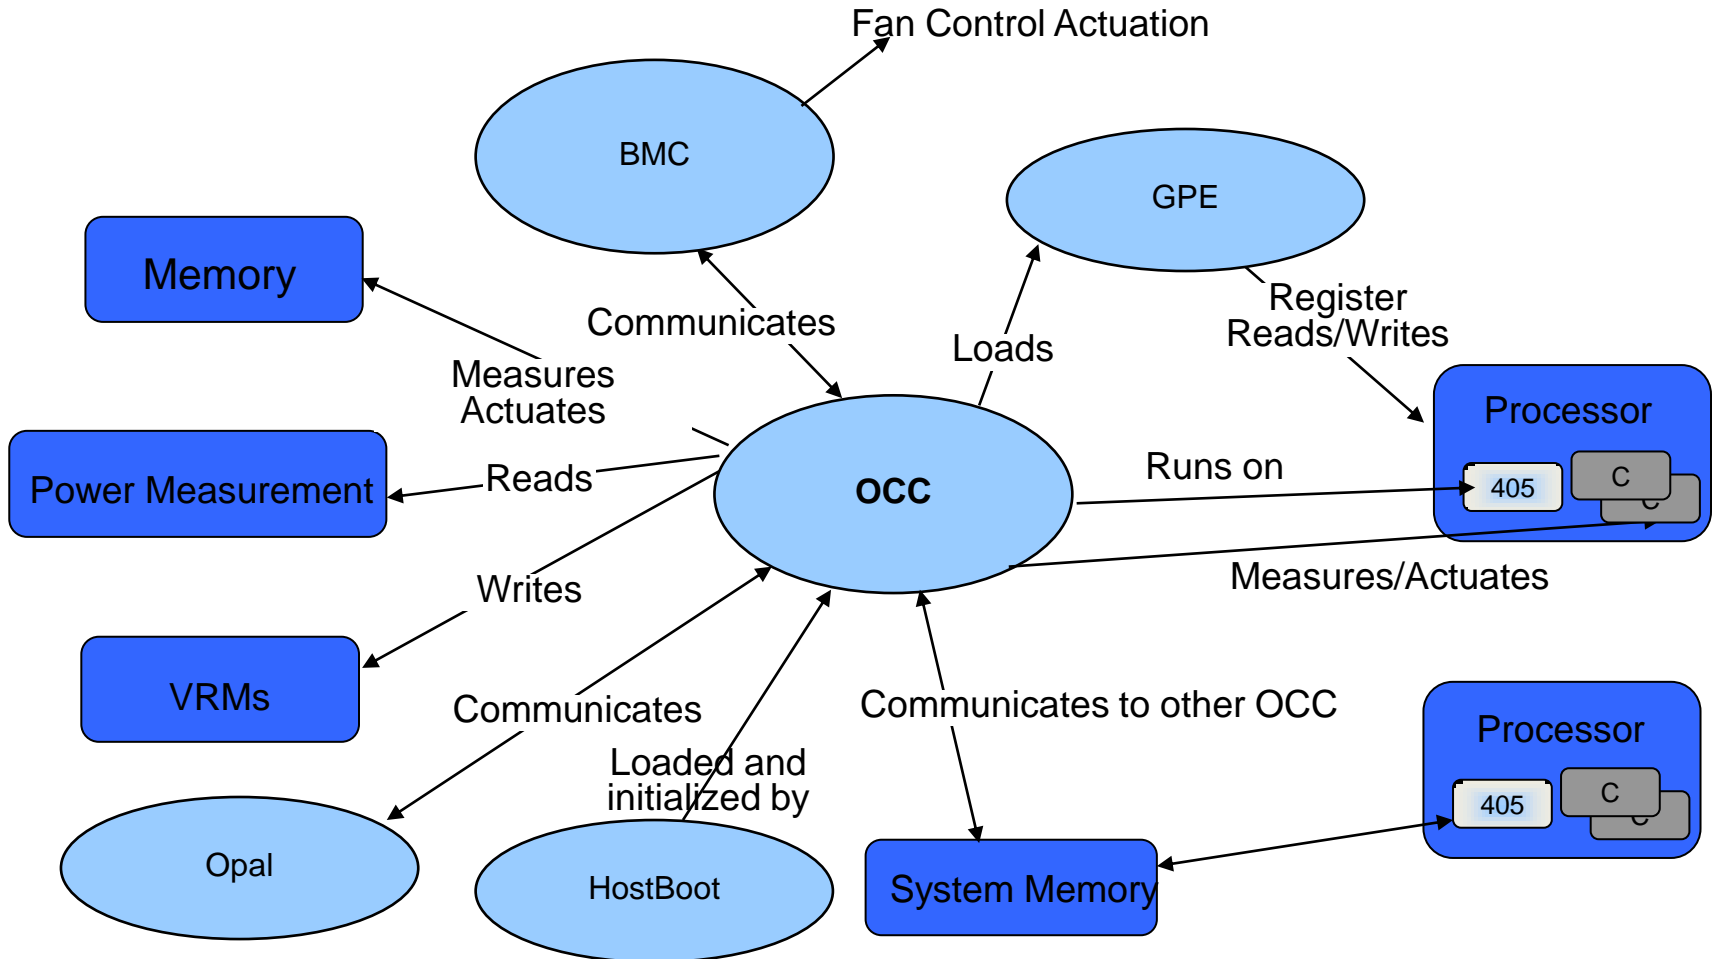

- Hardware/Firmware that controls power, performance & thermal
- 405 processor with 512k dedicated RAM
- General Purpose Engines (GPE) to offload the 405

What does OCC do?

- Reads/controls system power
 - Enables efficient fan control
 - Provides OT protection
- Power Capping
- Fault Tolerance
- Energy saving
- Performance boost



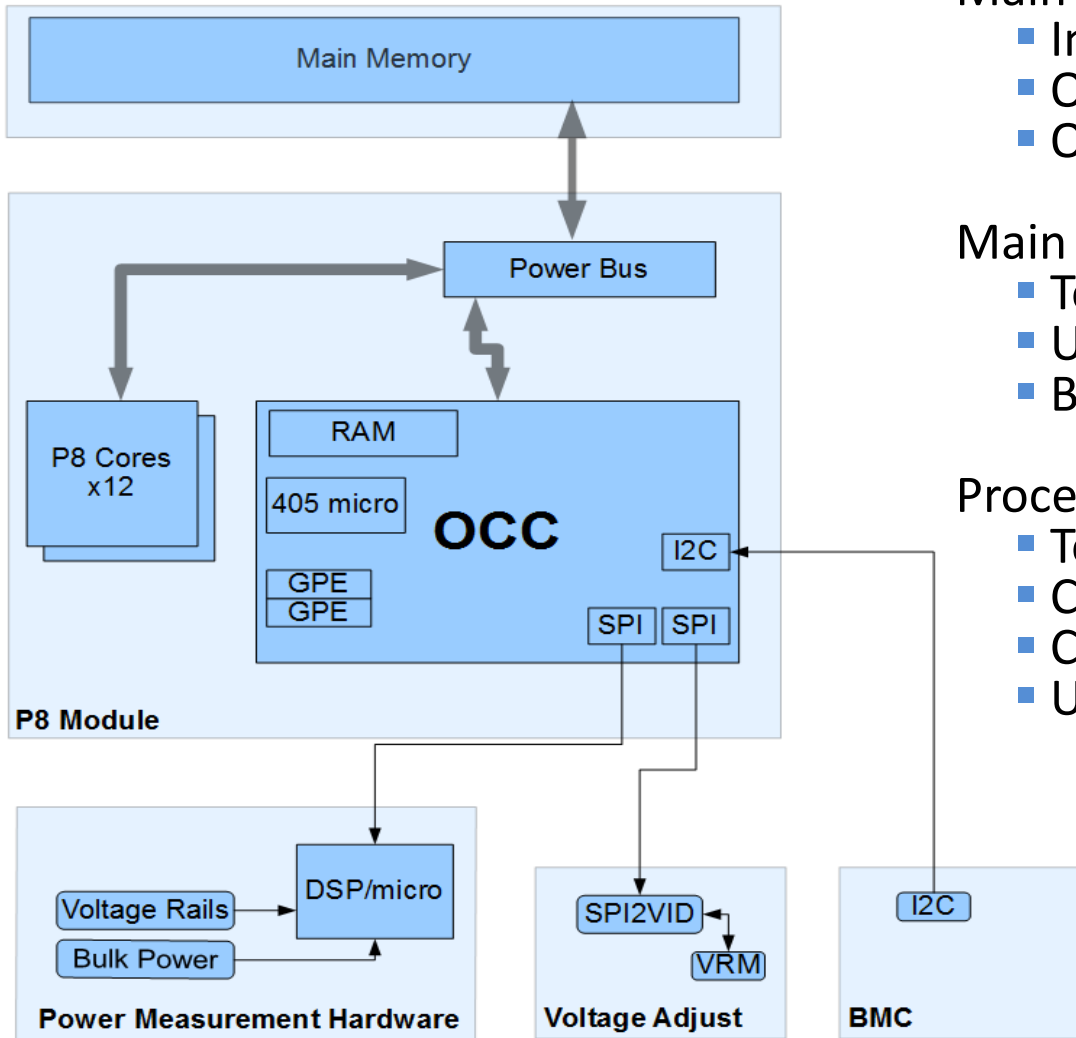

OCC Overview



Physical paths not shown



Hardware Block Diagram



Main Memory Uses

- Initial Load
- OPAL communication
- OCC communication

Main Memory Measurement

- Temperature Sensing
- Utilization Measurement
- Bandwidth control

Processor

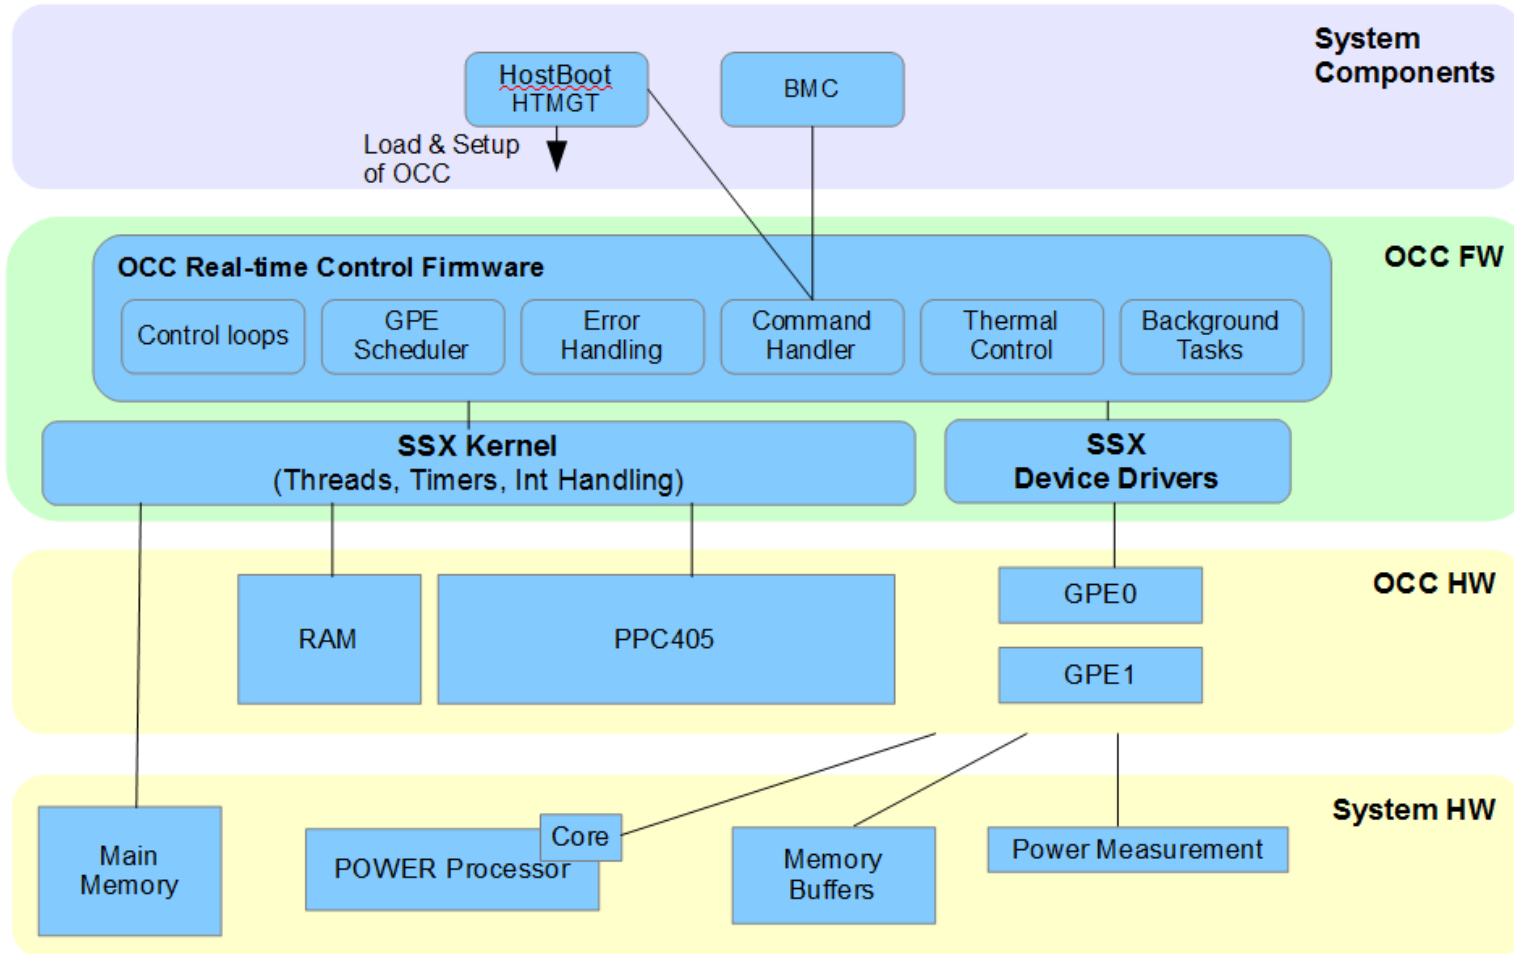
- Temperature Sensing
- Core Frequency Control - PSTATES
- Chip Voltage Control – PSTATES
- Utilization Measurement

BMC


- Report power/temperature
- Provide Power Cap
- DCMI compliance



OCC Internal Details




PSTATE Details




Utilization and governor policy are inputs to cpufreq in Hypervisor



Cpufreq – exploits SPR based P-state change



User Set Power Limit



Processor Over Temperature



■ Power Shifting Improvements

- Memory Bandwidth % vs Processor Frequency
- Power Shift between other devices

■ Modify Controls

- Run processors at higher speeds
- Run processor/memory at higher temperatures

■ Adjust loop times

- Increase responsiveness to workload changes

Open Power Blog link:

- <http://openpowerfoundation.org/press-releases/occ-firmware-code-is-now-open-source/>

GitHub pages:

- OpenPOWER project: <https://github.com/open-power>
- OCC: <https://github.com/open-power/occ>
- Building OpenPOWER: <https://github.com/open-power/op-build>

Technical Firmware Training – OpenPOWER Pavilion:

- Thursday 9-11am Architectural Overview
- Thursday 12-2pm Hands-on Training – Build/Modify

Reference Platform (Palmetto) – TYAN GN70-BP010