UEFI Power Calibration & Optimization Tools

SSG: Rahul Khanna, Daryl McDaniel
DCG: Christian Le
Intel Corporation

Agenda

- Objective
- Power Calibration & Optimization Technologies
- Romley Implementation
- MemLoad Tool & Demo
- Future Tools

Objective

Maximizing energy efficiency, performance per Watt, is paramount for future data centers. This requires technologies which can dynamically evaluate and limit power on platform and component level. We present an technology to calibrate and optimize power in platform.

1) UEFI Tools for Pro-Active Power Calibration Technology: Tools with algorithm to stress and calibrate weights in UEFI for Romley.
2) Accurate Power Estimation Algorithm utilizing memory performance counters and platform configuration specific weights provided by UEFI tool
3) Demonstrated on: Intel® Xeon® 5500 Reference Server

UEFI Tool for Pro-Active Power Calibration Technology

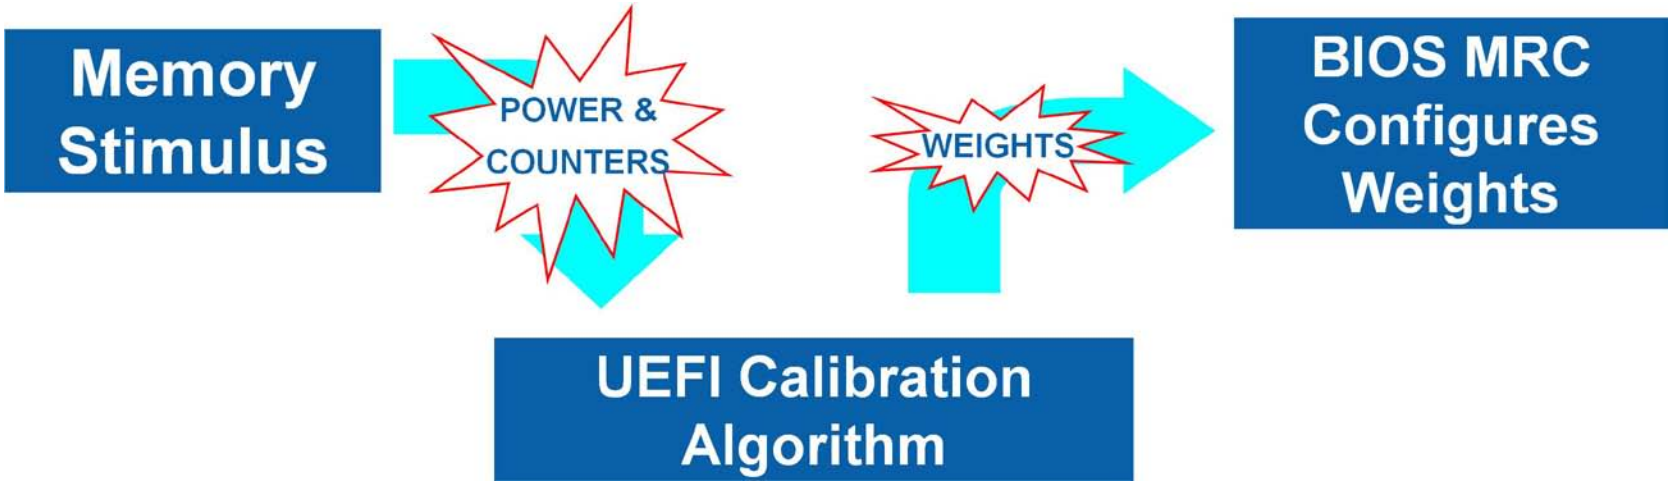

Implementation of UEFI Memory Calibration Tool

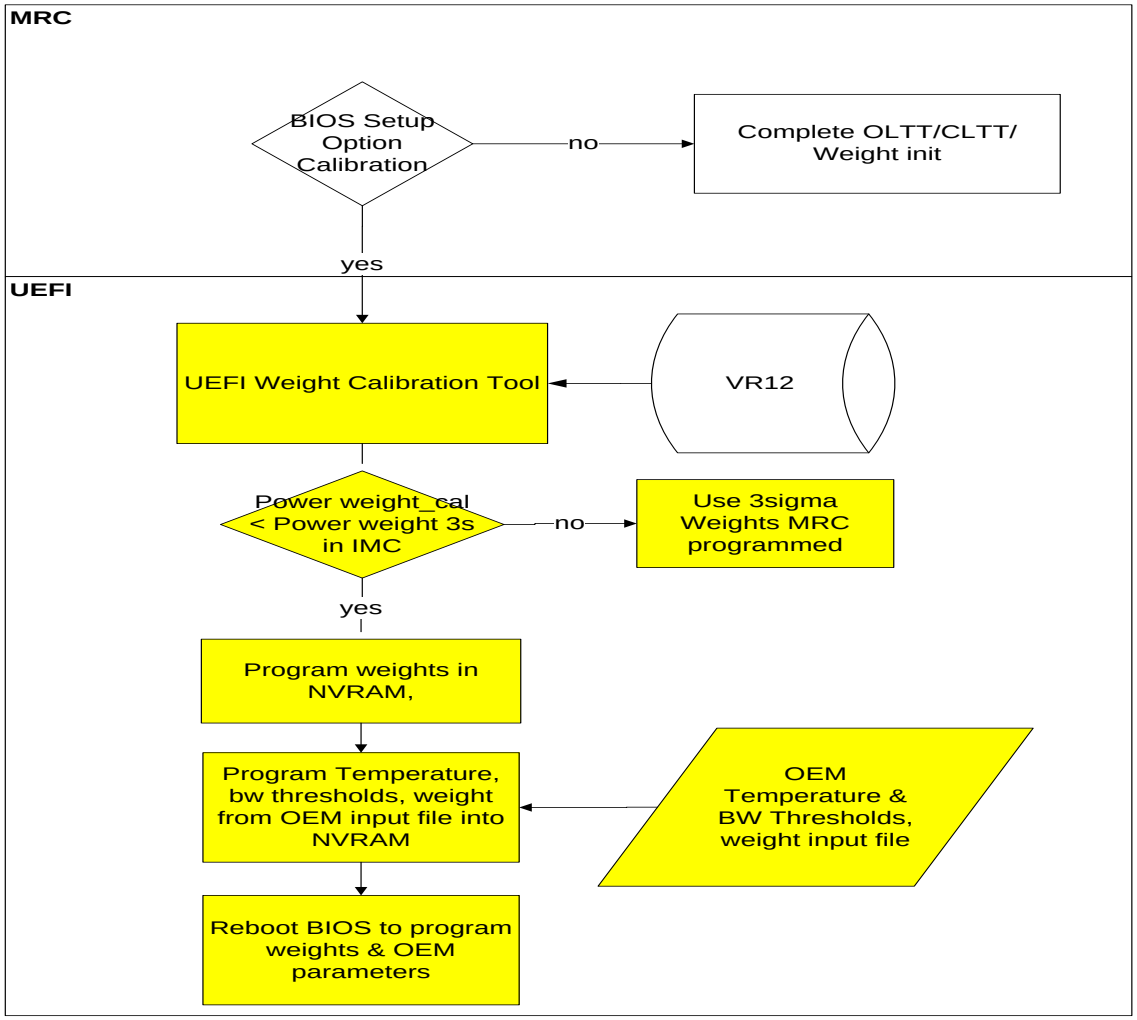

Complex Algorithms Migrate to Tools in EFI

MemLoad Features

- Native UEFI Application
  - No OS Overhead
  - Platform Independent
- Deterministic Loading
- Tunable Loads
- Configurable
  - Platform Configuration
  - Processor / Memory association
  - Memory Channel targets
  - Local / Remote NUMA loading
  - Command Line & Control File

- Multi-Processing
  - Uses UEFI specified APIs
  - One process per core/thread
  - Boot core provides control
  - Load processes on APs
  - Direct control over load process assignment to core/thread

MemLoad Internals

- One control process
  - Application Setup
  - Starts load processes
  - Gathers load statistics
  - Calculates access latencies
  - Displays results
- **N** load processes
  - Targets two memory regions
    - Local + Remote, or
    - Remote + Remote, or
    - Local + Local
  - Accesses to different cache lines
    - Ensures memory accesses
- Per-pass Statistics
  - Bytes transferred per second
    - Aggregate of all processes
  - Transaction access time
    - Meas. By the control process
- Tunable parameters, Per-pass
  - Pass duration, in seconds
  - Idle time per access, microseconds
  - Random or Sequential accesses
  - R/W or Read Only
  - Read/Write Ratio

Demonstration: Intel® Xeon® 5500 Reference Server

Accurate Memory Power Estimation

- Intel® Xeon® 5500 Reference Server w/ SW Drivers & Tools
- Average 15% lower power with calibration

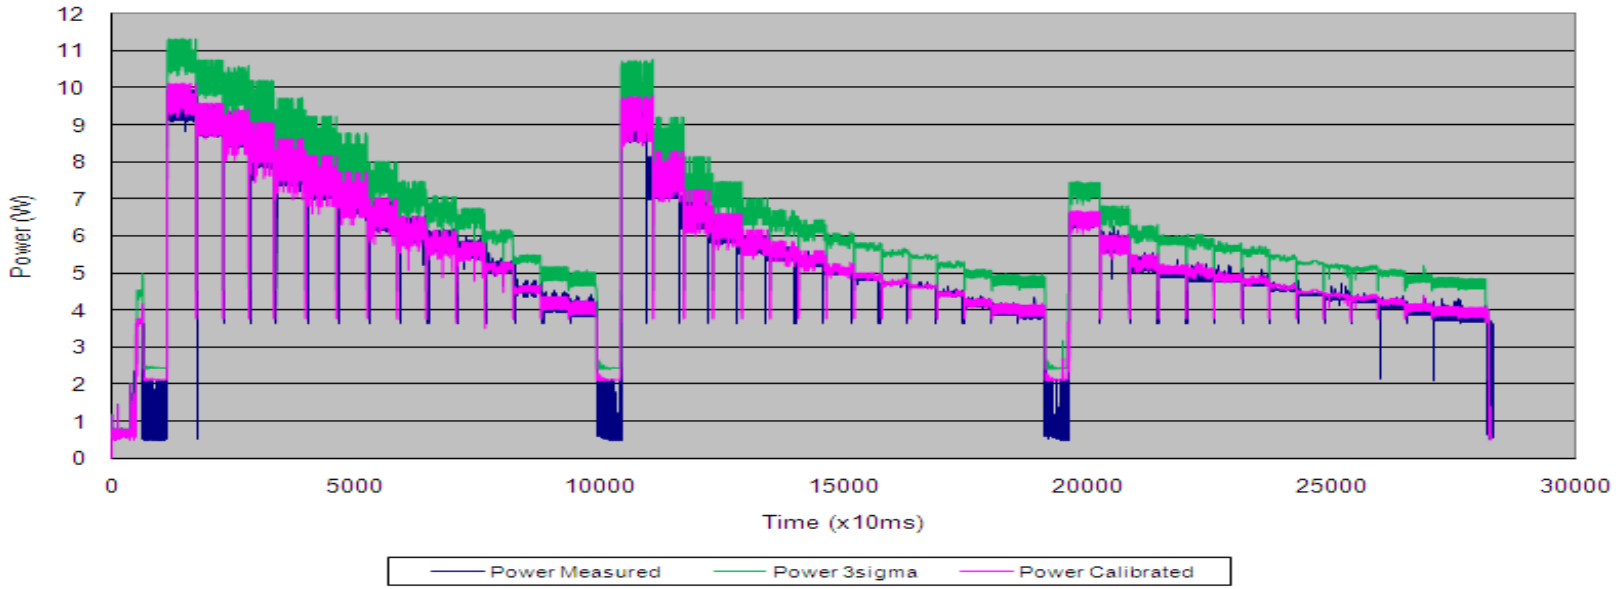

Next Steps

- Future UEFI tools CPU & platform level power calibration & optimization tools
- UEFI value-add Tools are licensed from Intel – Contact Fadi Zuhayri fadi.zuhayri@intel.com