RNET PowerAware: Fine Grained Power Monitoring for High Performance Computing Applications

RNET Technologies, Inc.

RNET's PowerAware platform provides fine-grained power monitoring capability to application developers. The platform performs high-frequency, fine-grained power and energy profiling for all system components. The profiling correlates the power and energy profiling to fine-grained application phases. This allows the application to monitor and tune its power and energy footprint on modern power constrained strained large-scale computing systems.

Power consumption is a primary challenge that must be addressed by HPC infrastructure and application developers on the path to exascale computing. RNET's PowerAware solution can help developers capture a fine-grained view of their application's power and energy consumption.

The PowerAware platform consists of the PowerAware monitoring board (a PCIe expansion card), component power monitoring sensors, host driver, and the user-level PowerAware API. The PowerAware board collects power and energy information from system components (e.g., CPU, memory, compute accelerators, network cards, disks, fans, etc.) via the sensors and generates statistics. Power data and statistics can be processed using multiple mechanisms; transferred to the host (via PCIe), transferred to a remote host via Ethernet, or buffered on PowerAware board's local storage for future retrieval.

The PowerAware host API can be used to measure power or energy consumption of any system component for arbitrary phases of application execution. The power/energy profiles can be used to highlight and further improve each phase's power and energy footprint. The application can also perform online power actuation to dynamically adjust power consumption of each phase.

Technical Specification

Sampling Frequency
Up to 32KHz

Analog Power Channels
32, expandable to 128

ADC Sampling Resolution
16 bits

Processing Unit
Xilinx Spartan6 FPGA

On-board Memory
256 MiB DDR2

On-board Storage Capability
SD and CompactFlash

Connectivity
1GE, PCIe x1

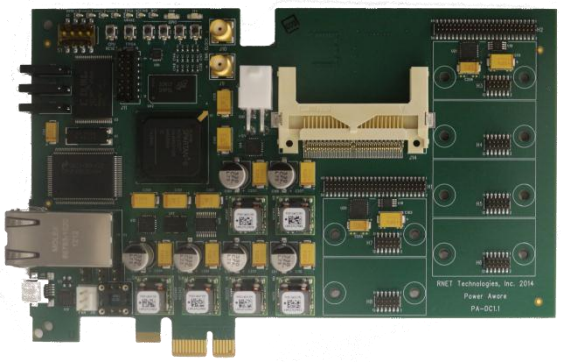

Fine Grained Power Monitoring

RNET PowerAware

The PowerAware Software Architecture and API

The PowerAware software architecture includes the board firmware, the kernel driver, and the host API. The PowerAware host API is the primary entry point for interacting with the PowerAware platform. The API provides access to each power monitoring board installed on the system. Future versions of the API will also be able to access remote PowerAware platforms connected via the Ethernet interface. Using this API the user can start power monitoring tasks, monitor power channels (i.e., system hardware components), and calculate power or energy for any subset of the power channels in arbitrary application phases specified by the application or profiling tool. The API also allows the user to set callbacks that can be executed when power/energy measurements meet user-defined thresholds.

The PowerAware Hardware Architecture

The current and voltage measurements are collected on the PowerAware board through small sensor boards and risers that are installed close to the system components. The PowerAware board contains 16-bit, 1 MHz ADCs that are used by the Xilinx Spartan FPGA to read the analog sensor output. The PowerAware board aggregates and processes the data based on the user requirements and synchronizes the sensor measurements with application execution to determine application phases. The processed and raw data can be stored on on-board storage (i.e., CompactFlash or SD cards), sent to the host via the PCIe interface, or sent to a remote host via the Ethernet interface.