Extending the Lifetime of SSD Controller

Author: Deepak Shankar
Tel : 408-569-1704
Fax : 408-519-6719
Email: dshankar@mirabilisdesign.com
Website : http://www.mirabilisdesign.com/

Abstract

- Developed performance models to evaluate the efficiency of SSD Controller
- Input was varying workloads and interfaces
- Looked at varying the wear leveling, data distribution across the flash devices and different garbage collections
- Generated reports around the effective bandwidth, Read/Write latency
- Compared the reports with the typical operating specification of the vendor

Motivation for Architecture Simulation

- Complex behavior
  - Input rate and task sequence
  - Data size and priority dependent behavior
- Contention
  - Limited resources on the platform
  - Scheduling/arbitration of task and diagnostics
- Interference of multiple applications
  - Competing for resources
  - Scheduling/arbitration
  - Unexpected anomalies

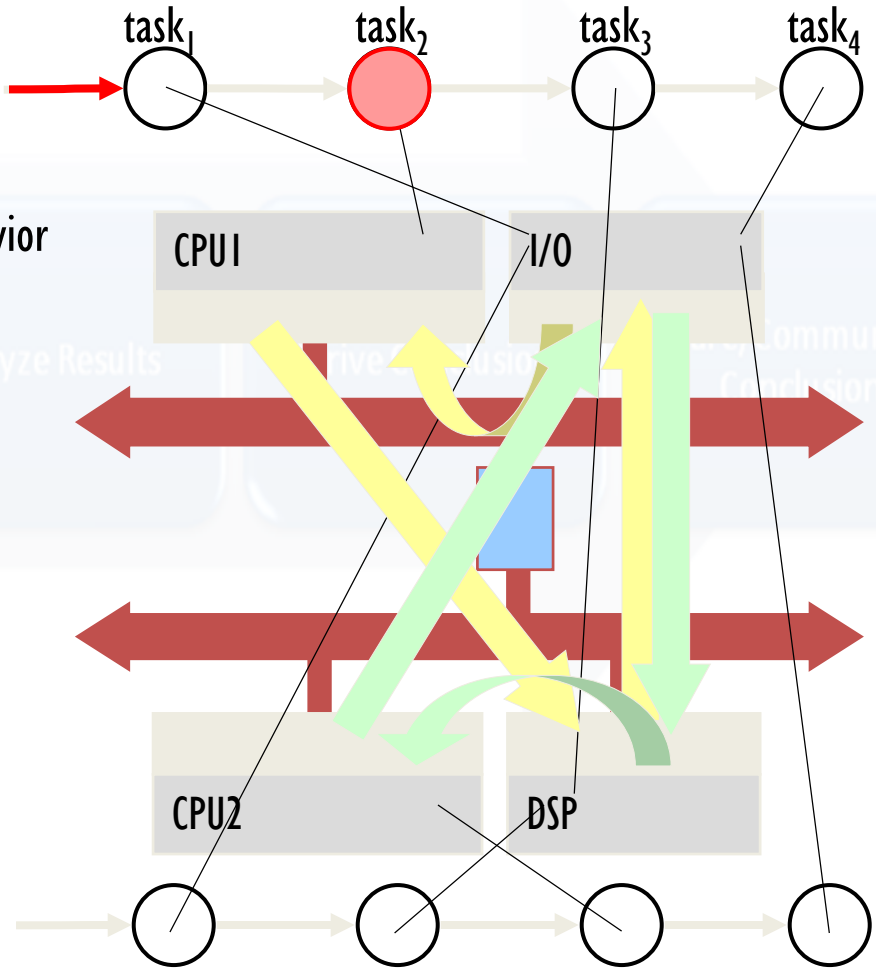

Justification for System-level Model

System with faster Bus is slower in places

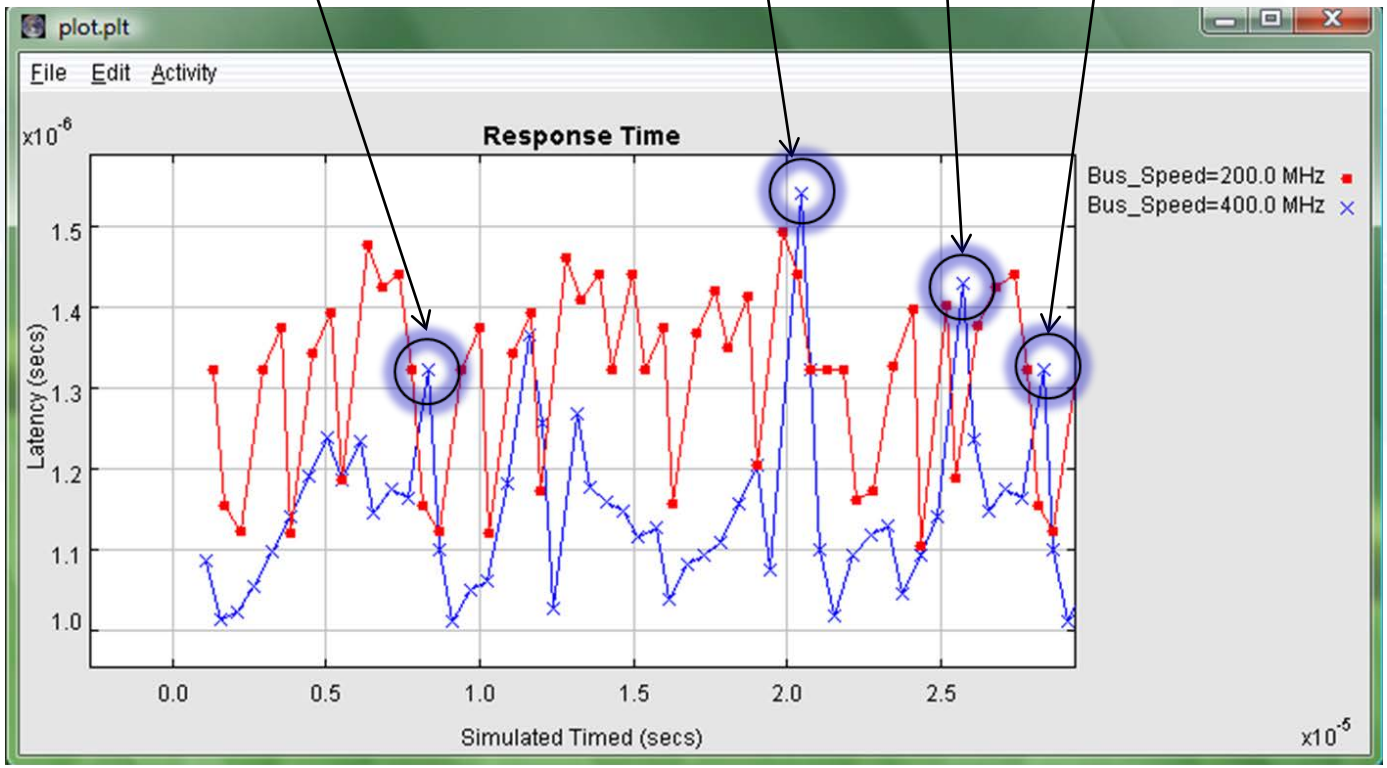

Unpredictable system response

Results

- Life of an SSD is directly dependent on the write amplification factor (WAF)
- WAF in turn is related to SSD overprovisioning, which is a parameter that the system designer can control
- A diligent system designer can extend the life of an SSD by upto 60% by proper control of over-provisioning, thus reducing Total Cost of Ownership (TCO).

Variations

- At $1-2 a 1GB of SSD,s, TCO of datacenters deploying SSDs has a huge dependency on the effective life of an SSD
- Measuring dynamic and simulated outputs of critical metrics such as WAF, over provisioning is critical is assessing the life of the SSD
- Future dynamic studies planned include
  - Separation of hot and cold data
  - File system stream management
  - Look at additional ways to extend the life of the SSD.

Methodology Adopted

- Used a Architecture Simulation environment
- Constructed a statistical workload with a functional description of the system
- Size of each SSD: 256GB
- Developed an array of 32 devices
- Connected PCIe + NVMe
- Using random distribution generator with varying request sizes, priority and rates
- Created a network of multi-processor and multi-core

Architecture Model of NVMe SSD

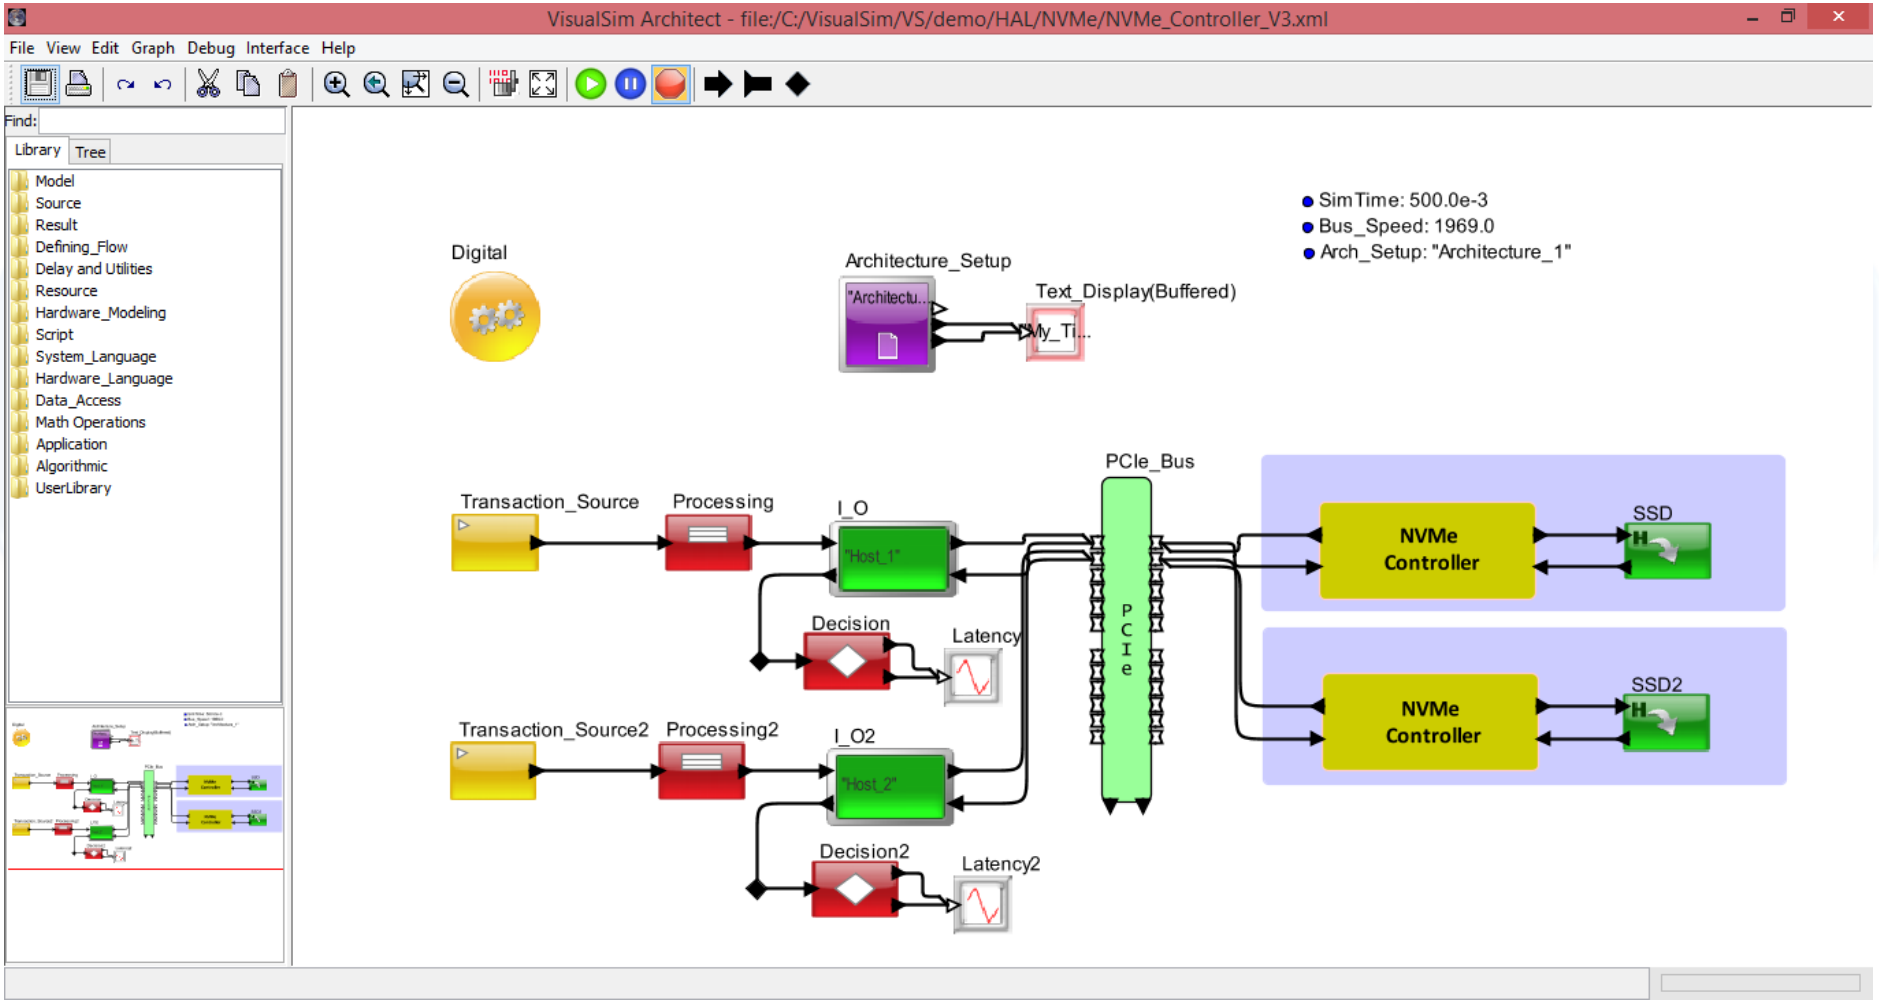

Output and Reports

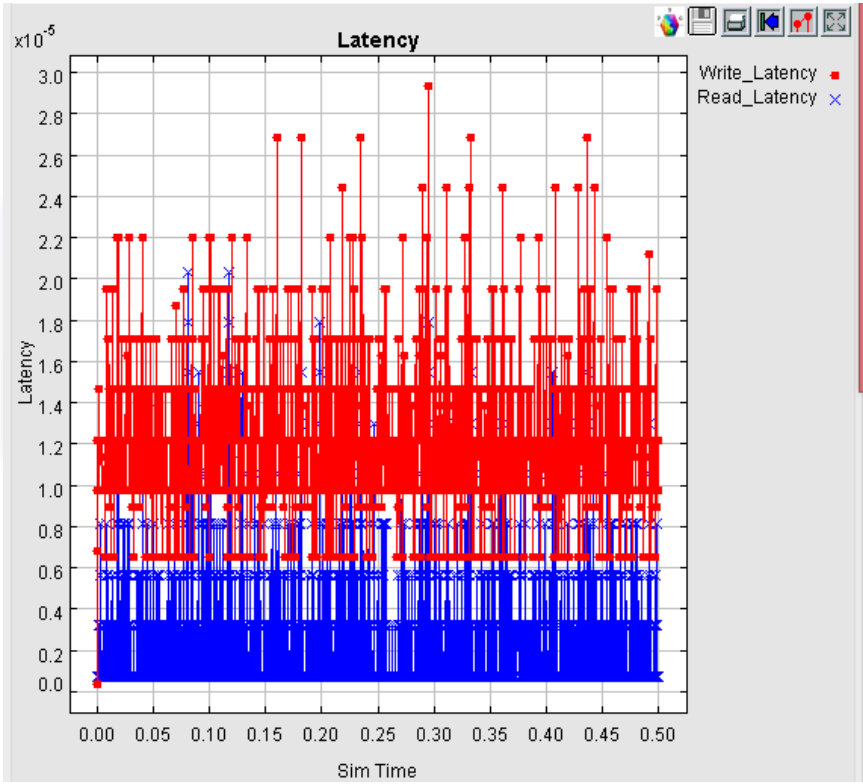

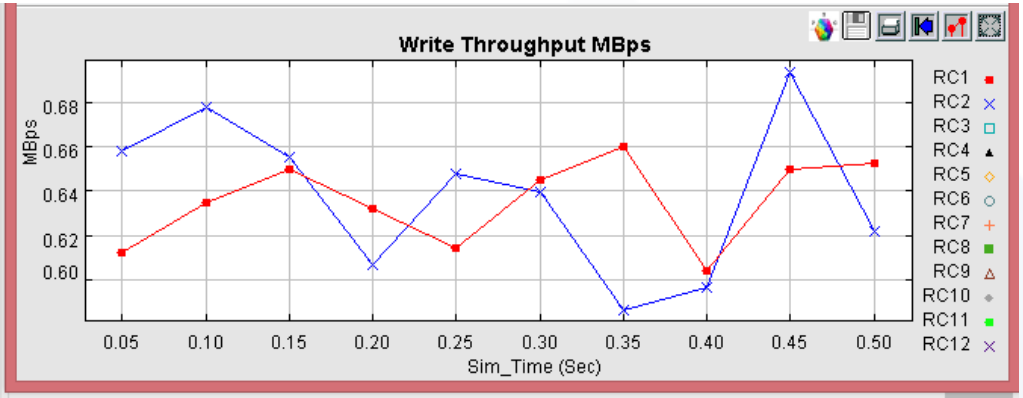

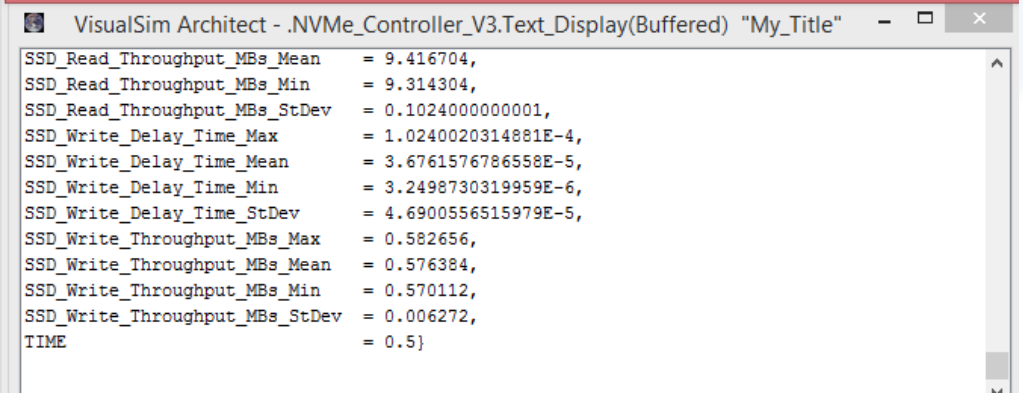

Focus Areas of Analysis

- Functionality
  - Quality, correctness and accuracy of flow
  - Define network/interface/bus protocols, arbitration, schedulers, mode selection, logic flows
  - Effectiveness of Diagnostics
- Performance
  - Latency, Throughput, utilization
  - Buffer, Processing, Bandwidth, Hit-miss, operational delays, burst vs. stream
  - Parallel processing, messaging, resource consumption
  - Devices, workload, traffic rate, behavior flows, threads

Focus Areas of Analysis (Cont.)

- Power
  - State based power definition for each device
  - High accuracy
  - Assign unique states for different operations
  - Incorporate leakage, static and dynamic power
  - Define power management logic as state machines
  - Power state changes at cycle-level
  - Dynamic power activity based on workload and device status
  - Linear approximation for capacitance and inductance
  - Standard plots for battery usage, instant power, average power, per device and per task
  - Can account for changes in temperature, clock speed and availability of charging resources

VisualSim- Modeling and Simulation

- Architecture exploration and system validation solutions
- Graphical modeling, exploration and analysis
- Custom, statistical and cycle-accurate modeling libraries
- 200+ application templates to accelerate development
- Pre-built blocks enabled for performance and power
- Over 15 interfaces and extensive algorithm definition library
  - C/C++, SystemC, Java, MatLab/Simulink, Verilog, Python etc.

400 building blocks, custom modeling functions and full system visualization

About Mirabilis Design

- Solution for product definition, communication and adoption
- Using system-level modeling and simulation
- To design Systems, FPGA, Processor/SoC and Real-Time Software
- With over 500 system modeling IP with timing and power
- And in-house experts in system modeling and analysis
- Having the largest number of electronic system design users

Select the "Right" configuration to match customer request

MIRABILIS design

System Simulation and Exploration

Corporate Headquarters

1159 Sonora Ct, Suite 116

Sunnyvale, CA 94086

USA

Tel : 408-844-3234

Fax : 408-519-6719

Sales Information : info@mirabilisdesign.com

Technical Contact : tech.support@mirabilisdesign.com

Website : http://www.mirabilisdesign.com/

http://www.visualsim.com