

Sandia's HPC Hardware Architecture Strategy

Development of our foundation for Exascale

August 30, 2011

James A. Ang, Ph.D.

**Scalable Computer Architectures Department
Computing Research Center
Sandia National Laboratories
Albuquerque, NM**

HPC Paradigm: Custom versus Commodity

- The last paradigm shift in HPC was the move from Cray vector supercomputers to massively parallel processor (MPP) supercomputers
- This revolutionary change was known as:
The Attack of the Killer Micros - Eugene Brooks, LLNL
 - Founded on a philosophy of leveraging the rapid advances that were available from commodity microprocessors that rode the wave of both Moore's Law and Dennard Scaling
 - MPPs based on commodity microprocessors killed Cray Research's custom vector supercomputer business
- The ASCI Program established critical mass for this paradigm shift by investing heavily and equally in:
 - MPP application development
 - Computer science and enabling technologies
 - Large scale platforms

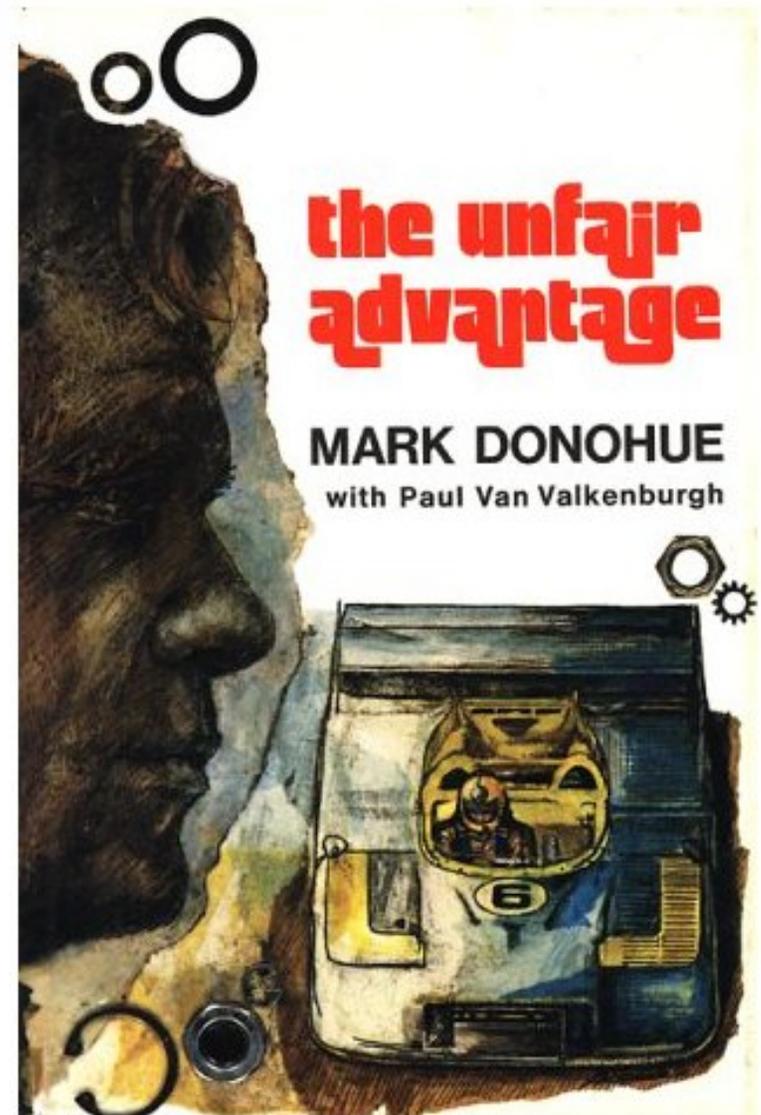


The HPC Paradigm is Primed for the next Major Change

- **Five years ago commodity microprocessors began to change**
 - Dual core processors appeared due to power and cooling limits, and commodity processors began to fall short of performance needs for HPC
 - Multi/many core exacerbate the memory wall/data movement problem
 - The result is a growing performance gap between theoretical and realized performance for our real applications
- **Co-design is an implicit statement that multi-core commodity processors need to be redesigned with an eye towards the needs of HPC**
- **The assumption is improvements in these new processor designs will still be mainstream so the HPC community can benefit from re-designed commodity processor volumes**
 - Sandia can play a critical role in *Crossing the Chasm . . .*

The Unfair Advantage

- Early in the development of the 917, Mark Donohue proved to Porsche that the Penske team he drove for was not like the other race teams
- The *Unfair Advantage* he enjoyed was based on his ability to communicate with Porsche engineers on their terms
 - More than a race car driver, Donohue was also a Mechanical Engineer
 - Donohue was directly involved in the development of the Porsche 917: suspension design, turbocharger sizing, fuel injection tuning, etc.
- In 1975, Donohue drove the Can-Am Porsche 917/30 to a lap speed of 221.16 MPH at *Talladega* – a closed course record that stood for over 10 years
- Donohue was a *Lead User*





The Issue / Our Challenge:

Commodity processor adoption of capabilities for HPC

- How are HPC co-design innovations integrated into commodity processor designs?
- The MPP HPC paradigm, while based on commodity processor designs, has never influenced those designs
- HPC may now have an opportunity because Industry has more transistors than they know what to do with
 - Stamping out more cores that will be even more starved for data is an indication that Industry may be receptive to good ideas
 - Perhaps especially good ideas from *Lead Users* who have an *Unfair Advantage*

Key Focus Areas

**Help Develop and Exploit Sandia's Unfair Advantage:
Leverage our capabilities as a collaborator with Industry
– not just a customer**

- **Develop Strategic Capabilities for Co-Design**
 - SST: Open Framework for HPC Architectural Simulators
 - Mantevo: Proxy Apps to drive SST and use on existing platforms
 - Objective Function Design: Not just minimizing *Time to Solution*, must also balance with minimizing *Energy to Solution*
- **Data Movement Research and Development**
 - Interconnection Networks: Smart NICs, Portals, Si Photonics, etc.
 - Advanced Memory: Stacked DRAM, integrated NVRAM, etc
 - *Anyone can build a fast CPU. The trick is to build a fast system.*
– Seymour Cray
- **Help develop and support ASC Platform Strategy**
 - Cielo, ACES 2015 Platform
- **Help develop and support DOE Exascale Initiative**

Key Focus Areas - Continued

- **Cultivate Sandia's Unfair Advantage**
 - **Micron Cooperative Agreement**
 - **PM support for X-caliber, Mgmt support for XGC**
 - **Procure Experimental Architecture Testbeds**
 - Intel MIC
 - AMD Fusion
 - Convey HC-1ex
 - Tiler TILE Gx-36
 - **Support NNSA/ASC Program Office (Starting Sept'11 @ 50% time)**
 - New ASC Platform Strategy – CD0 for ACES 2015 Platform
 - Re-establish PathForward Program
 - Help Thuc Hoang draft new Computing Plan / CSSE Tech. Roadmap
 - **Support DOE EI: E7 Exec, SPEC Exec, ACES Arch Office**
- **Develop HW Prototypes w/ Sandia's μ electronics Center**
 - **Exascale Grand Challenge LDRD**
 - **Leverage X-caliber proposal, collaboration on Quilt Packaging**

Co-design for HPC

- **Lessons from the embedded computing community**
 - Working with Prof. Sharon Hu, University of Notre Dame, one of the pioneers of co-design for embedded computing design
 - Optimization is based on partitioning the work among different elements of the embedded system to minimize energy consumption
- **Ongoing and New Co-design efforts**
 - DARPA/UHPC X-caliber project
 - Sandia Exascale Grand Challenge LDRD project
 - DOE/ASCR Data Movement Dominates project
 - ACES – Cray Advanced Interconnection Network project
 - Participant in two of the recently announced DOE/ASCR Exascale Co-design Centers:
 - Combustion ECDC
 - ExMatEx CDC

Industry Collaborations

Cultivate strategic partnerships

- Intel – Umbrella CRADA
- IBM – Trusted Foundry Work, informatics
- Cray – CRADA for informatics & ACES D&E project
- Oracle (Sun) – Red Sky
- Micron Technology – CRADA for ECC development

How do we encourage industry to pursue the revolutionary approaches needed to address Exascale?

- Growing expertise in computer engineering and architecture – important to collaborate with industry as partners
- This expertise compliments our application, algorithm, and system software developers to be able to co-design and develop prototypes to implement and experiment with revolutionary models of computation
- We are also making a sustained investment in the open development of the SST simulation framework, and Mantevo mini-application proxies
- With our Microsystems design and fabrication capability, Sandia is also able to create proof of concept prototypes and hardware artifacts

Summary of Micron-Sandia Interactions

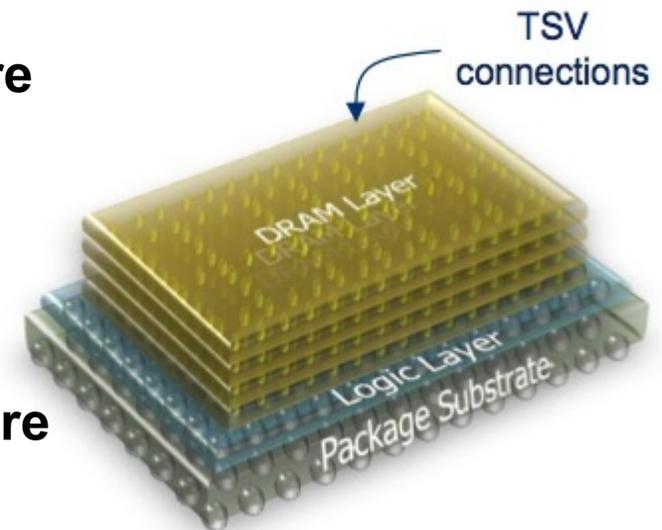
Activity	Outputs
Kickoff Meeting	Micron-Sandia Collaboration Begins (July'06)
Foundation for Discussions	2-way NDA (Nov'07)
Advanced Memory for DOE Architectures	Simulations, Papers, PIM LDRD Effort, ASC/CSSE L2 Milestone: Evaluate Advanced Memory Subsystems – 4Q-FY10
IAA Activities	Dean Klein is a member of the IAA Advisory Board IAA Workshop on Memory Opportunities for HPC (Jan'08) IAA Workshop on HPC Architectural Simulation (Sept'09)
Collaborations with other agencies	Alignment of ASC/CSSE, DoD/ACS, & IAA support to Integrate U-MD's Memory Simulator (DRAMsim & eBOBsim) with Sandia's SST
Proposal Partnerships	DARPA/UHPC X-caliber, DOE/ASCR Data Movement Dominates
CRADA – established July'10	Micron-Sandia collaboration to analyze advanced concepts for error correction in advanced memory designs – Patent Application filed, May '11, <i>Automated discovery of optimal, symbol-based SECDED codes</i>

- **Technical Exchanges from July 2006 - Present**

- Approximately 30 face-to-face technical meetings
- Catalyst for Collaboration with other DOE/NNSA labs, DoD, and Universities

New Micron Technology Inc. Cooperative Agreement with NNSA/ASC

- Based on our 5+ year collaboration, Sandia is helping HQ establish a D&E project & Cooperative Agreement with Micron
- On behalf of NNSA/ASC and DOE/ASCR, Sandia is responsible for technical oversight of and collaboration with Micron
- Sandia working with Micron to define the technical scope:
 - Construction of the simulation infrastructure
 - Design explorations, focused on sets of in-memory operations to be evaluated, simulated, and analyzed
 - FPGA prototyping work of ideas that are identified as candidates for inclusion in future Hybrid Memory Cube (HMC) parts
 - Research into improved low energy signaling and topologies



*Functioning prototypes in silicon
TODAY*

Microsystems and Engineering Sciences Applications (MESA) Fabrication Facilities

Silicon Fabrication Facility



- Total clean room area 33,000 ft² (12,500 ft² Class 1)
- 5V Bulk and 3.3V SOI Rad Hard CMOS in production
- In-house microelectronics technology & facility to deliver specialized IC products
- Primary supplier of custom Rad-Hard ICs for weapon life extension programs and Satellite Systems
- Supports silicon bulk and silicon surface micromachining
- DOD Defense Microelectronics Activity (DMEA) accredited Microelectronics Trusted Supplier (design and foundry services)

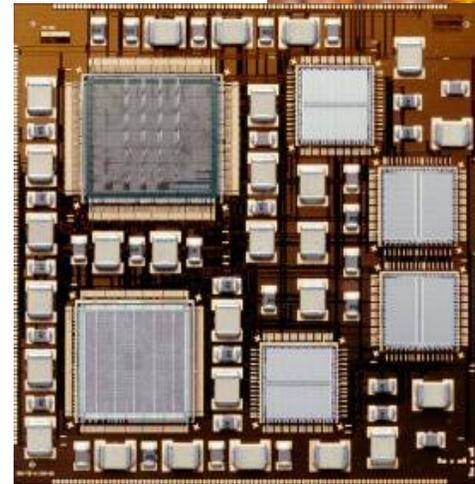
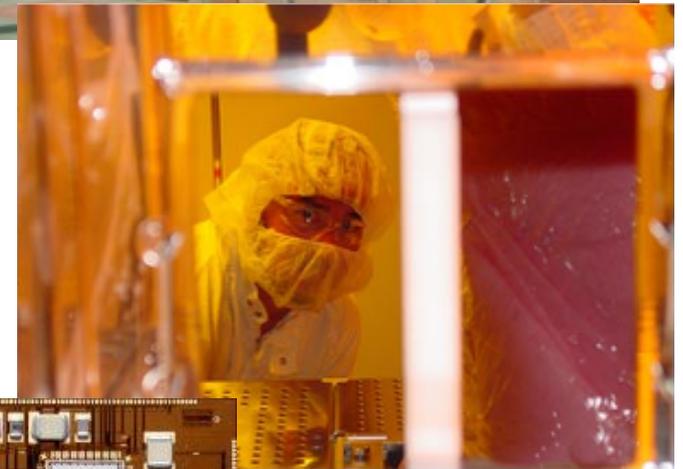
Microfabrication Facility



- 89,000 ft² (16,640 ft² Class 10/100)
- Reconfigurable tools from wafer pieces to 6" wafers
- 6" silicon post-processing facility to support hybrid substrates and 3D integration (8" compatible)
- Compound Semicond. Epitaxial Growth
- Compound Semicond. Discretes, IC's and MEMs
- Si MEMS / photonics / optoelectronics / VCSELs
- Mixed-Technology Integration and Processing
- 3D Integration; Packaging
- Materials Characterization
- Failure Analysis

Leverage MESA Capabilities for Exascale

- **Key areas of collaboration:**
 - Processor/Memory design
 - 3D Integration and Quilt Packaging
 - Optical Interconnects:
 - VCSELs
 - Si Photonics
- **Sandia's Microsystems Center is one of the heaviest users of IBM's trusted foundry**
- **Recent trusted foundry usage**
 - FY06-08: 14 Design Submissions (130nm, 90nm nodes)
 - FY09-10: 16 Design Submissions (130nm, 90nm, 65nm, 45nm nodes)



Some Keys to Realizing Exascale Computing

- **Improving data movement performance / efficiency**
 - Both intra-node and inter-node
- **Developing a Co-design methodology for HPC**
 - Including the development of the objective function that balances minimizing time to solution with minimizing energy to solution
- **Providing processor, memory, and interconnection network designers with insights from application users/ developers, system software developers on the best ways to use the additional transistors that Moore's Law will continue to provide**
- **Significant investment in new application development**
 - Both funding and time