SAND2013-6462P

Exceptional service in the national interest

Sandia National Laboratories

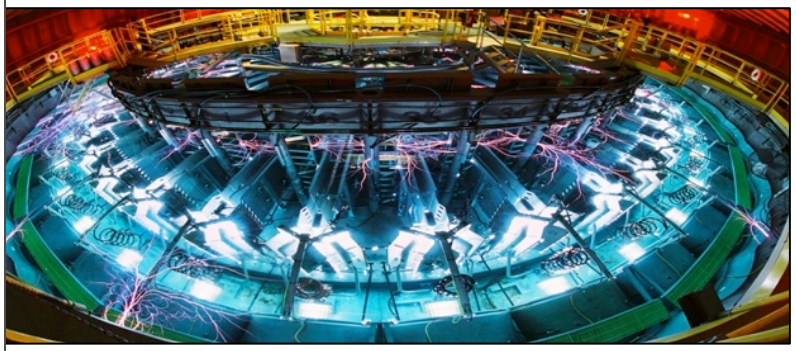

NNSA/ASC Test Bed Update

S.D. Hammond, R.F. Barrett, C.T. Vaughan. J. Jayaraj, C.R. Trott, J.H. Laros, S.M. Kelly, J.A. Ang
and System Admins

U.S. DEPARTMENT OF ENERGY | NNSA National Nuclear Security Administration

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Methodology for Codesign

Sandia National Laboratories

- Key Co-design Capabilities
- HPC Architectural Simulators
- Proxy Applications
- Advanced Architecture Testbeds
- Proxy Architectures

- Mantevo Proxy Applications
- http://www.mantevo.org
- SST HPC Simulation Framework
- http://code.google.com/p/sst-simulator

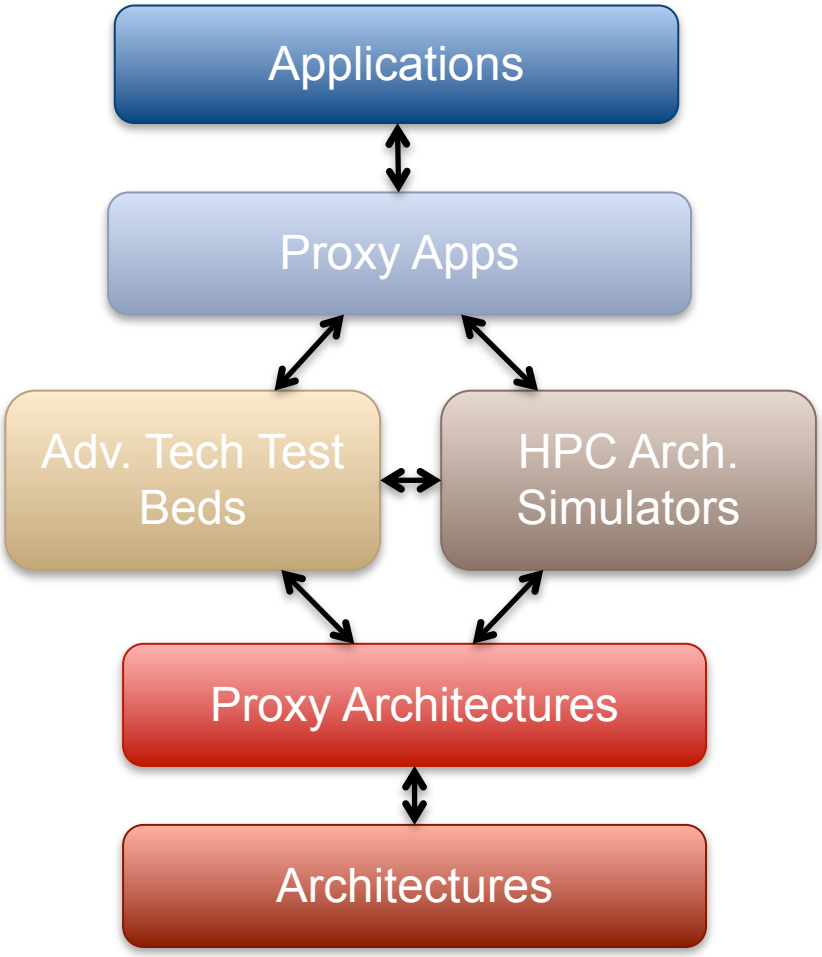

NNSA/ASC Test Beds

Sandia National Laboratories

- **Compton**: Dual-socket Intel Sandy Bridge + Dual Xeon Phi co-processor with InfiniBand interconnect
- **Shannon**: Dual-socket Intel Sandy Bridge + Dual NVIDIA Kepler K20X GPUs with InfiniBand interconnect
- **Teller**: AMD A10 Fusion "Trinity" (Quad-core + HD 7XXX series Northern Islands GPU) with InfiniBand interconnect
- **Watson**: Dual-socket IBM POWER7+ with 10GbE

- **Cray XC30**
- **Others**

Compton Test Bed

Sandia National Laboratories

- Compton is a test bed cluster for Intel's Xeon Phi product

- 42 nodes
- Dual-socket, 8-core 2.66GHz Sandy Bridge E-class processors
  - HPC processors found in TLCC-class machines, SMT is enabled on some nodes
- Each processor has a PCIe link to a Xeon Phi (Knights Corner) card
  - Knights Corner is a 57-core, 4-way SMT thread card
  - 5.0GT/s memory bus (up to 120 – 140GB/s of STREAM B/W)
  - 8GB of GDDR5 for each card
- Each node has an 80GB SSD to enable NFS relationship between node and Xeon Phi cards
- InfiniBand interconnect for nodes (and Xeon Phi cards)

Shannon Test Bed

Sandia National Laboratories

- Shannon is a test bed cluster for NVIDIA's Kepler K20-GPU

- 32 nodes
- Dual-socket, 8-core 2.66GHz Sandy Bridge E-class processors
  - HPC processors found in TLCC-class machines, SMT is enabled on some nodes
- Each processor has a PCIe link to an NVIDIA K20X "Kepler" GPU
  - 2688 "Cores" @ 732MHz
  - 6GB GDDR5 Memory
  - 1.3 TFLOP/s
- InfiniBand interconnect for nodes

Teller Test Bed

Sandia National Laboratories

- Teller is a test bed for AMD's APU (CPU+GPU) product

- 104 nodes
- Single-socket, A-10, quad-x86-core APU, 3.8GHz
  - Desktop-grade part being used to evaluate potential server-grade parts available in the future
- GPU on each APU is integrated
  - Radeon HD 7660D (Northern Islands)
  - Is AMD's older GPU VLIW architecture
- Each node has 16GB of system RAM with 2GB separated for the GPU aperture
- "Zero-copy" between CPU and GPU is enabled at the driver level
  - Experimental-grade feature which works but is not perfect
- InfiniBand interconnect

Snapshot of Technologies

Sandia National Laboratories

	MPI	Intrins.	CUDA	OpenAcc	OpenCL	Kokkos Array	Cilk+	OpenMP		TBB	ArBB/ CEAN	qthreads	pthreads	MKL/ Math Lib.	Adv. Lang.	DSLs
								S	T							
NVIDIA (GPU)																?
AMD (CPU)																?
AMD (APU)		?														?
Intel (CPU)																?
Intel (MIC)																?
IBM (BG)										?						?
ARM																?
miniFE																
miniMD																
miniGhost																
LULESH																
S3D																
CoMD																

Successes

Sandia National Laboratories

- Many application ports now completed
  - OpenMP, OpenACC, OpenCL, CUDA, MPI
- Demonstrating significant performance improvements on some hardware (for some problems)
  - 2x not uncommon, 4x in some cases
- Growing experience in how to port and optimize code
- Excellent vendor discussions and feedback
  - Compilers, vendor libraries, profilers, debuggers etc

Difficulties

Sandia National Laboratories

- Technology is still emerging
  - OpenACC is particularly problematic neither CAPS or PGI really work (even compile) with C++
  - Many, many bugs in the software environment
- Tension between depth of code changes which we want to pursue
  - Deeper, intrusive changes are giving better performance
  - Application developers are saying they don't want to over specialize code
- Keeping stack up to date is challenging
  - We are building many components ourselves

Observations

Sandia National Laboratories

- Vectorization is proving challenging
  - And not *always* beneficial (particularly in solvers)
- Very often optimizations are making codes on Xeon run faster
  - Can recognize some performance improvements even today
- Focusing on parallelism seems to result in fluid algorithm design
  - We keep thinking "MPI + Threads + Vectors"
  - In the end most platforms blend different amounts of these
- Use MPI between NUMA domains, threads inside
- Don't use #pragma …

Sandia
National
Laboratories

Exceptional service in the national interest