



# GPU Clusters at NCSA

**Guochun Shi**

**Innovative Systems Laboratory  
(ISL)**

National Center for Supercomputing Applications  
University of Illinois at Urbana-Champaign



# Innovative Systems Laboratory (ISL)

- **Work with technology vendors to evaluate new computing technologies for possible future use in HPC; Technologies of interest include:**
  - Computational accelerators (FPGA, Cell, GPU, ...)
  - HPC architectures (e.g., clusters with accelerators)
  - Interconnects
  - HPC system software (e.g., Microsoft Windows cluster solution)
- **Work with computational scientists on the evaluation of the suitability of new compute technologies for applications and the amenability of existing codes for new computing technologies. Example applications include:**
  - Computational chemistry, cosmology, quantum chromodynamics, etc.

# GPU Clusters at NCSA

- **Lincoln**

- Production system available the standard NCSA/TeraGrid HPC allocation
- nodes: 192
- CPU cores: 1536
- Accelerator Units (S1070): 96
- Total GPUs: 384
- CPU cores/GPU ratio: 4

- **AC**

- Experimental system managed by NCSA staff. The GPUs are donated by NVIDIA to Wen-mei Hwu's research team.
- nodes: 32
- CPU cores: 128
- Accelerator Units (S1070): 32
- Total GPUs: 128
- CPU cores/ GPU ratio: 1




# AC Cluster




# AC: Compilers and Cluster Management

- **Programming tools**
  - CUDA C 2.2 SDK
    - CUDA/MPI
  - OpenCL 1.0 SDK
  - PGI+GPU compiler
- **Cuda\_wrapper**
  - Intercepting a subset of CUDA/openCL APIs
  - Enable sharing of one node
  - Virtualization of physical GPUs
  - Ensure NUMA affinity to GPUs
- **Cuda\_memtest**
  - Ensuring integrity of the GPU memory
  - Keep nodes in healthy state

# TeraChem



Bovine pancreatic  
trypsin inhibitor (BPTI)  
3-21G, 875 atoms, 4893  
basis functions



## MPI timings and scalability