

IDF: Multi-Core Processing for HPC

March 2005




Russ Miller

Center for Computational Research

Computer Science & Engineering

SUNY-Buffalo

Hauptman-Woodward Medical Inst



University at Buffalo

The State University of New York

21st Century University

- **Embrace digital data-driven society**
- **Empower students to compete in knowledge-based economy**
- **Support research, scholarship, education, and community outreach**
- **Deliver *high-end cyberinfrastructure* to enable efficient**
 - **Collection of data**
 - **Management/Organization of data**
 - **Analysis of data**
 - **Visualization of data**

Center for Computational Research 1999-2005 Snapshot

■ High-End Computing, Storage, Networking, and Visualization

□ ~100 Research Groups in 37 Depts

○ Physical Sciences

○ Life Sciences

○ Engineering

○ Scientific Visualization, Medical Imaging, Virtual Reality

□ 13 Local Companies

□ 10 Local Institutions

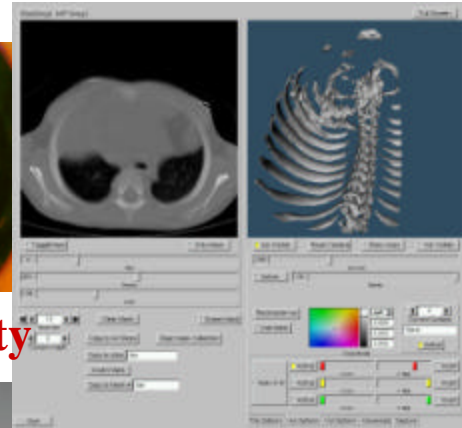


■ External Funding: \$300M+

■ Total Leveraged WNY: \$0.5B

■ Deliverables

□ 1100+ Publications

□ Software, Media, Algorithms, Consulting,
Training, CPU Cycles...




Major Compute/Storage Resources

- **Dell Linux Cluster (2.9TF)**
 - ❑ 600 P4 Processors (2.4 GHz)
 - ❑ 600 GB RAM; 40 TB Disk; Myrinet
- **Dell Linux Cluster (6TF)**
 - ❑ 4036 Processors (PIII 1.2 GHz)
 - ❑ 2TB RAM; 160TB Disk; 16TB SAN
- **IBM BladeCenter Cluster (3TF)**
 - ❑ 532 P4 Processors (2.8 GHz)
 - ❑ 5TB SAN
- **SGI Intel Linux Cluster (0.1TF)**
 - ❑ 150 PIII Processors (1 GHz)
 - ❑ Myrinet
- **RFP (10-15TF)**
 - ❑ Pentium-Based
 - ❑ Fast Interconnect
 - ❑ Efficient Storage Management
- **SGI Altix3700 (0.4TF)**
 - ❑ 64 Processors (1.3GHz ITF2)
 - ❑ 256 GB RAM
 - ❑ 2.5 TB Disk
- **Apex Bioinformatics System**
 - ❑ Sun V880 (3), Sun 6800
 - ❑ Sun 280R (2)
 - ❑ Intel PIIIs
 - ❑ Sun 3960: 7 TB Disk Storage
- **HP/Compaq SAN**
 - ❑ 75 TB Disk; 190 TB Tape
 - ❑ 64 Alpha Processors (400 MHz)
 - ❑ 32 GB RAM; 400 GB Disk

CCR Visualization Resources

- **Fakespace ImmersaDesk R2**
 - ❑ Portable 3D Device
 - ❑ Onyx2: 6 R10000 @ 250MHz
 - ❑ 2 IR2 Pipes; 3 64MB texture memory mgrs.
- **Tiled-Display Wall**
 - ❑ 20 NEC projectors: 15.7M pixels
 - ❑ Screen is 11' ´ 7'
 - ❑ Dell PCs with Myrinet2000
- **Access Grid Nodes (2)**
 - ❑ Group-to-Group Communication
 - ❑ Commodity components
- **SGI Reality Center 3300W**
 - ❑ Dual Barco's on 8' ´ 4' screen
 - ❑ Onyx300: 10 R14000 @ 500MHz
 - ❑ 2 IR4 Pipes; 1 GB texture mem per pipe



Multi-Core Applications

- **Application Performance is Key**
 - “Moore’s Law of Applications”
- **Processor (Socket/Core) Performance Irrelevant**
- **Balanced System is Critical**
 - Memory and Cache Access including SMP
 - System I/O
 - *No Choke Point*
- **Multiprocessing is Critical**
 - Human Ingenuity Required/Available to Decompose Algorithms
 - Rededication to the Shared Memory Programming Problem
- **Multithreading is Important**
 - Difficult Problem in Scientific Computing
 - Tools Required for Multi-Core/Multithreading Environment

Intel Multi-Core Systems

- **Pentium Dual Core – 2005**
 - Desktop Environment; Serves as Development Platform
- **Itanium Dual Core (Montecito) – 2005**
 - Multithreading \neq Dual Socket System = 8 Processors to OS
 - 12MB L3 Cache
 - Healthy Thermals
- **Xeon Dual Core (Dempsey) – 2006**
 - Blackford/Greenfield Chipset – Balanced System
 - Dual Independent FSB
 - FB-DIMMs More Efficient for Large Memory Configurations
- **Tukwila – 2007/8**
 - More than 2 cores
 - Common Chipset and Commodity Components (Memory, Power Supplies, etc) Aimed at Consolidating Xeon and IPF

Benefits of Intel Multi-Core Systems

- **Application-Based Moore's Law requires**
 - **Parallel (Distributed and/or Shared Memory) *and***
 - **Multithreaded Implementations of Algorithms**
- **Tools for Multithreading are Critical**
 - **Hand-Coded Libraries (BLAS, SHMEM, FFT)**
 - **Posix Threads**
 - **Compiler Directives (OpenMP)**
 - **Hybrid MPI-OpenMP**
- **Balanced Systems are Critical**
- ***Well-funded applications will reap the benefit***

Multi-Core Application Development

- **More Emphasis on Computational Science & Engineering at all Levels**
- **Black Box Approach is Counter Productive**
- **Programming Skills Must be Improved**
- **Training Required at all Levels**
 - High School
 - Undergraduate
 - Graduate
 - Post-Doctoral
- **Better & More Affordable Tools are Required**
 - Intel Provides Compilers and Libraries Free of Charge (Academic Use/No Support)
 - Intel is Major Player in Development Software (Costly)

