Summarizing the September HPC User Forum meeting

Some Highlights

Steve Finn
steve@sfinn.com
Cherokee Information Services Inc.
Chairman, HPC User Forum Steering Committee

Meeting Topics

- HPC in Materials Science * (Ping Golf)
- Datacenter Power and Cooling * (TACC only)
- The path to Exascale * (Sriram Swaminarayan only)
- Clouds for HPC ?
- The Future of Lustre *
- Invited vendor presentations
- Other invited speakers
  - 1000 Genome Project
  - The AFRL Condor Cluster
  - Update on HECToR
  - Data Intensive Computing at SDSC *
  - Networking Information Technology Research and Development (NITRD)
  - HPC as an Ecosystem
  - IDC Updates

Final agenda: http://www.hpcuserforum.com/registration/sandiego/sandiegoagenda.pdf

Videos of some talks: http://insidehpc.com/category/events/hpc-user-forum/ *talks marked with * available*

Presentations: http://www.hpcuserforum.com/download.html

Materials Science

- Presentations from:
  - Tim Germann, Los Alamos National Laboratory
  - Bobby Sumpter, Oak Ridge National Laboratory
  - Charles R. Welch, US Army Corps of Engineers
  - Andres Jaramillo-Botero, Caltech
  - Eric Morales. Ping Golf
- Common theme among 1^{st} 4 speakers on need to match codes to future architectures including challenges of heterogeneous and hierarchal aspects of future architectures

Tim Germann

Co-design is a process by which computer science, applied math, and domain science experts work together to enable scientific discovery

- Hardware is changing dramatically
  - *Increased concurrency*
  - *Increased heterogeneity*
  - *Reduced memory per core*
  - *"Business as usual" is not going to work*
- Algorithms and methods will have to be rethought / revisited
  - *Flops are (almost always) free*
  - *Memory is at a premium*
  - *Power is a constraint for large scale systems*
  - *Resiliency is a challenge*
- Few domain scientists have the extended expertise "from hardware to application" to enable applications to run at exascale
- Success on the next generation of machines will require extensive collaboration between domain scientists, applied mathematicians, computer scientists, and hardware manufacturers

Tim Germann

Summary

- Our objective is to establish the interrelationship between algorithms, system software, and hardware required to develop a multiphysics exascale simulation framework for modeling materials subjected to extreme mechanical and radiation environments.
- This effort is focused in four areas:
  - *Scale-bridging algorithms*
    - » UQ-driven adaptive physics refinement
  - *Programming models*
    - » Task-based MPMD approaches to leverage concurrency and heterogeneity at exascale while enabling fault tolerance
  - *Proxy applications*
    - » Communicate the application workload to the hardware architects and system software developers, and used in performance models/simulators/emulators
  - *Co-design analysis and optimization*
    - » Optimization of algorithms and architectures for performance, memory and data movement, power, and resiliency

ExMatEx

Datacenter Power and Cooling Panel

- Tommy Minyard, Texas Advanced Computing Center
  - Energy efficient strategies in a large data center
  - Servers suspended in mineral oil
- Henry Tufo, University of Colorado Boulder
  - Modular datacenter built by Epsilon in Ontario Canada, assembled on-site , measured PUE of 1.1
- Steve Jones, Stanford University
  - Retrofitting an existing building as a datacenter
  - 560 dual socket nodes, 120 NVIDIA M2050 GPU’s
  - Chilled Water, PUE 1.5

Exascale

- Sririam Swaminarayan, Los Alamos National Laboratory
  - Focused on changes to programming models
- Peter Beckman, Argonne National Laboratory
  - International Exascale Software Project
  - Need for investment in software to handle future architectures, recent HPC focus has been on hardware but software needs keep up
- Jeff Nichols, Oak Ridge National Laboratory
  - Partnerships will be crucial for the success of Exascale initiatives, including with industry. Many think Exascale will cost us $5B over the next 5 years, or about $600M a year.

The AFRL Condor Cluster

Integration, Development and Results of the 500 TeraFlop Heterogeneous Cluster (*Condor*)

September 2011

Mark Barnell
Air Force Research Laboratory

Integrity ★ Service ★ Excellence

AFRL

The *Condor* Cluster

FY10 DHPI Key design considerations: **Price/performance** & **Performance/Watt**

84 head nodes
- **6 gateway access points**
- **78 compute nodes**
  - Intel Xeon X5650 dual-socket hexa-core
  - (2) NVIDIA Tesla GPGPUs
    - 49 nodes – (98) C2050
    - 15 nodes – (30) C2070
    - 14 nodes – (28) C1060
  - 24 GB RAM (*48GB)

1716 SONY Playstation3s
- STI Cell Broadband Engine
  - PowerPC PPE
  - 6 SPEs
  - 256 MB RAM

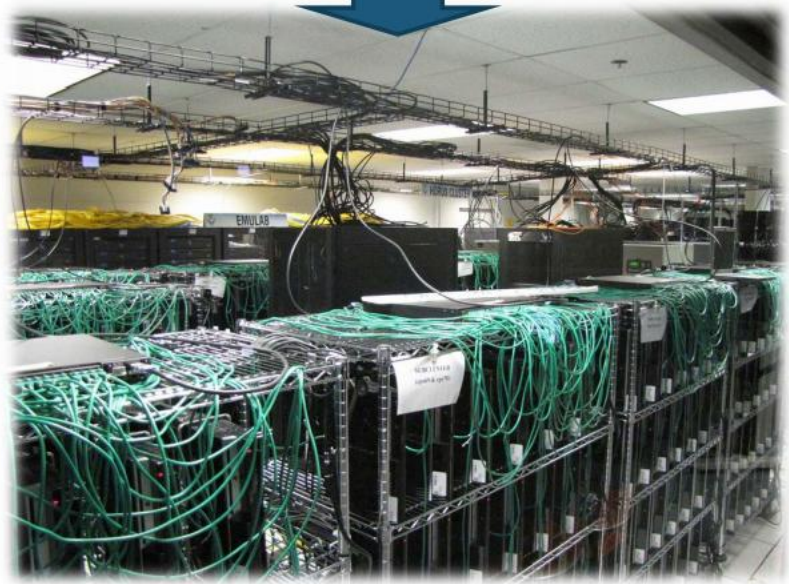

Solving Demanding, Real-Time Military Problems

Radar processing for high resolution images

Occluded text recognition

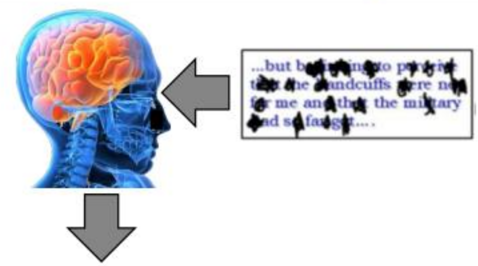

...but beginning to perceive that the handcuffs were not for me and that the military had so far got....

Space object identification

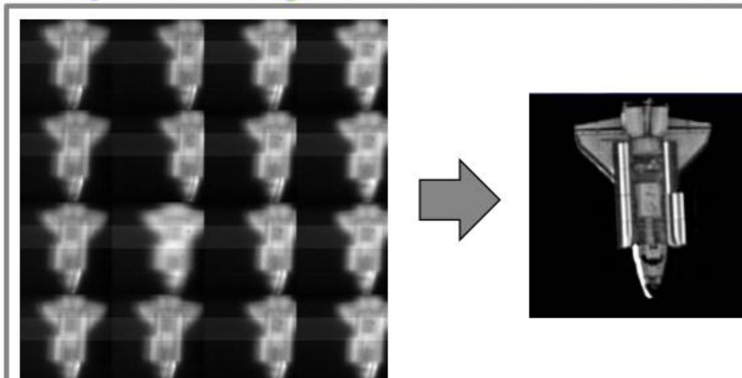

Skip Garner VBI

Analysis of the 1,000 Genomes data is enabling us to understand the basal level of variation in microsatellite loci – to discover new diagnostic markers, drug targets and toxicology tests

HPC Users Forum
September 7, 2011

Virginia Bioinformatics Institute
Virginia Tech

We have established a pipeline for the 1000 Genome Project and TCGA data

- Repeat 2,000,000 times per genome
- Thousands of genomes
- Data mine the finished product

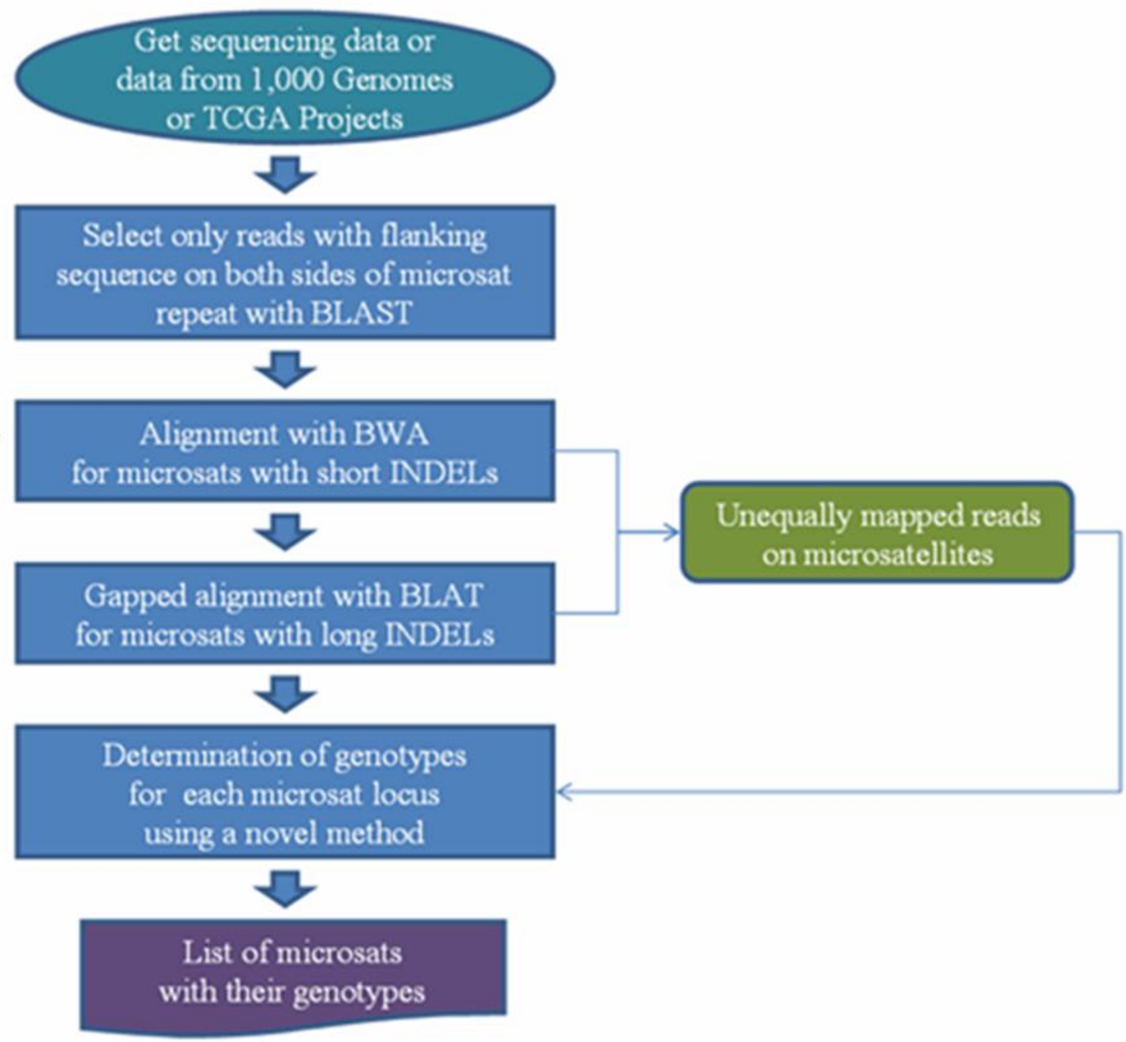

bwa aln part: ~4GB file (14 million 76 bp ready) takes 2 minutes on Convey HC-1. Or ~4 hours running on a sngle node 2x AMD Opteron 4174 (6 cores each, 2.8GHz, 6M Cache), 48GB RAM 1333MHz, with 4 NVidia Tesla GPU cards.