

Project Ideas Brainstorming

Ioan Raicu

**Center for Ultra-scale Computing and Information Security
Department of Electrical Engineering & Computer Science
Northwestern University**

EECS 395 / EECS 495

**Hot Topics in Distributed Systems: Data-Intensive Computing
January 19th, 2010**

Developing a research proposal

- Identify a problem
- Review approaches to the problem
- Propose a novel approach to the problem
- Define, design, prototype an implementation to evaluate your approach
 - Could be a real system, simulation and/or theoretical
- Write a technical report
- Present your results
- Write a workshop/conference paper (optional)

Project Ideas

- Tunebot on the Cloud
- Extending Erlang
- Distributed Data Mining in Weka/RapidMiner
- OS on Clouds/Distributed Systems
- Cross-origin AJAX
- Virtualization Impact for Data-intensive Computing
- Data aware scheduling on erasure codes based distributed file systems
- Grid type freedom in clouds and a more efficient utilization of resources
- QoS aware resource management in clouds
- Remote Function Passing
- Automatic parallelism discovery
- GPU implementation of computation-intensive tasks
- Data-intensive File systems
- Data Staging in Data-intensive computing
- Data Virtualization Service
- Data-Intensive Support for Manycore Architectures

Tunebot on the Cloud

- Tunebot is an online music search engine
- Queries take several seconds with a small database of 2500 songs
- What happens with many concurrent users, and a larger database? → major slowdowns
- Solution:
 - Parallelize
 - Port from a dedicated environment to a cloud one
- Evaluate scalability and performance

OS on Clouds/Distributed Systems

- Distributed Operating Systems
- Achieve a unified OS across machine boundaries
- The opposite of virtualization, which creates multiple virtual OS instances on one machine
- Choose an OS to modify
 - CPU scheduler → load balancing
 - Memory manager → shared memory
 - File system → leverage shared/parallel file systems
- Choose a virtual machine to modify (e.g. Java)
- Evaluate workloads for performance and scalability

Virtualization Impact for Data-intensive Computing

- Virtualization has overheads
- Quantify these overheads for a variety of workloads
 - Computational intensive
 - Memory intensive
 - Storage intensive
 - Network intensive
 - Across different virtualization technologies
 - Across different hardware
- Survey the latest research in addressing shortcomings of virtualization

Data aware scheduling on erasure codes based distributed file systems

- Distributed file systems use replication to ensure reliability of data
- Replication
 - Pros: Easy to implement, increases data locality and perf
 - Cons: Expensive and inefficient, in terms of network bandwidth and disk space
- Erasure codes:
 - Pros: Efficient in disk space usage
 - Cons: Harder to implement, expensive computationally, decreases locality
- Investigate replacing replication with erasure codes

QoS aware resource management in clouds

- The SLAs with cloud providers are limited
- Virtualization allows fine-grained control of physical resources
- Define ways to control network topology in Clouds
 - Modify an existing open source cloud middleware to add new APIs
- Propose pricing scheme for different levels of QoS
 - Evaluate pricing scheme; at what pricepoint do cloud resources have to be to match the QoS of grids?

Remote Function Passing

- Goal:
 - Maximize data locality in applications data access patterns
- Approach:
 - Move application to data
- Potential problems:
 - Load balancing
- Potential solutions:
 - Move data to application sometimes
- Involves data-aware scheduling algorithms and analysis

Automatic parallelism discovery

- Most code is inherently sequential in nature → this was OK while we doubled processor speeds according to Moore's Law
- Multi-core and manycore architectures are making sequential codes inefficient
- How to parallelize existing codes without burdening the programmer

HPC and Scientific Application on Manycore Architectures

- 100~1000 cores per GPU
- Does cluster computing programming approaches apply to GPUs?
- How can GPUs be generalized for HPC use?
- What architecture support is needed?
 - Cores should have L1/L2 caches, and GPU memory should be a L3 cache for the host memory → Nvidia Fermi might be a step in the right direction
 - Allow cores to execute independent kernels
 - No enforcement of coherency across cores
 - Allow core-to-core communication

Data-Intensive File Systems

- Implement a distributed file system
 - Use of FUSE for a general POSIX interface
 - Use structured distributed hash tables for distributed meta-data management
 - Can scale logarithmically with system size
 - Can create network topology aware overlays
- Relaxed data access semantic to increase scalability
 - eventual consistency on data modifications
 - write-once read-many data access patterns
- Evaluation scalability and performance
 - Compare to NFS, GPFS, PVFS, Lustre, HDFS

Data Staging in Data-Intensive Computing

- Most compute nodes are in 1 or 3 states
 - Input, compute, output
- Blocking I/O can yield low processor utilization
- There is a need for transparent mechanisms to overlap I/O with computations
- Project involves working and possibly modifying with HPC and MTC middleware (e.g. MPI, Swift, Falkon)

Virtual Replicas in HPC Systems

- High failure rate in modern HPC systems
 - Large number of components
 - Use of off-the-shelf unreliable components
- Failure rates dynamically varies based on
 - System architecture and Workload
- Replication for fault detection (possible tolerance)
- Independent virtual machines as replicas instead of stand-alone nodes

Questions

