Facilitate Parallel Computation Using Kepler Workflow System on Virtual Resources

Jianwu Wang[1], **Prakashan Korambath**[2], **Ilkay Altintas**[1]

[1] *San Diego Supercomputer Center, UCSD*

[2] *Institute for Digital Research and Education, UCLA*

Outline

- Kepler Scientific Workflow System
- Data-Parallel Scientific Workflow Scheduling via Kepler on Virtual Compute Cluster Resources

Mission of Scientific Workflow Systems

- Promote "scientific discovery" by providing tools and methods to generate larger, automated "scientific process"
- Support workflow design, execution, sharing, reuse and provenance
- Design efficient ways to connect to the existing data and integrate heterogeneous data from multiple resources
- …

Kepler Scientific Workflow System

http://www.kepler-project.org

- Kepler is a cross-project collaboration: over 20 diverse projects and multiple disciplines.
- Open-source project; latest release available from the website
- Builds upon the open-source Ptolemy II framework
- Vergil is the GUI, but Kepler also runs in non-GUI and batch modes.

… initiated August 2003

- 1.0 release: May 13th, 2008
- 2.0 release: June 24th, 2010
- 2.1 release: Sep 30th, 2010
  - *More than 40 thousand downloads!*

Ptolemy II: A laboratory for investigating design
KEPLER: A problem-solving support environment for Scientific Workflow development, execution, maintenance

KEPLER = "Ptolemy II + X" for **Scientific Workflows**

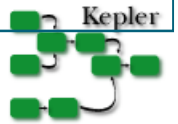

Actors are the Processing Components

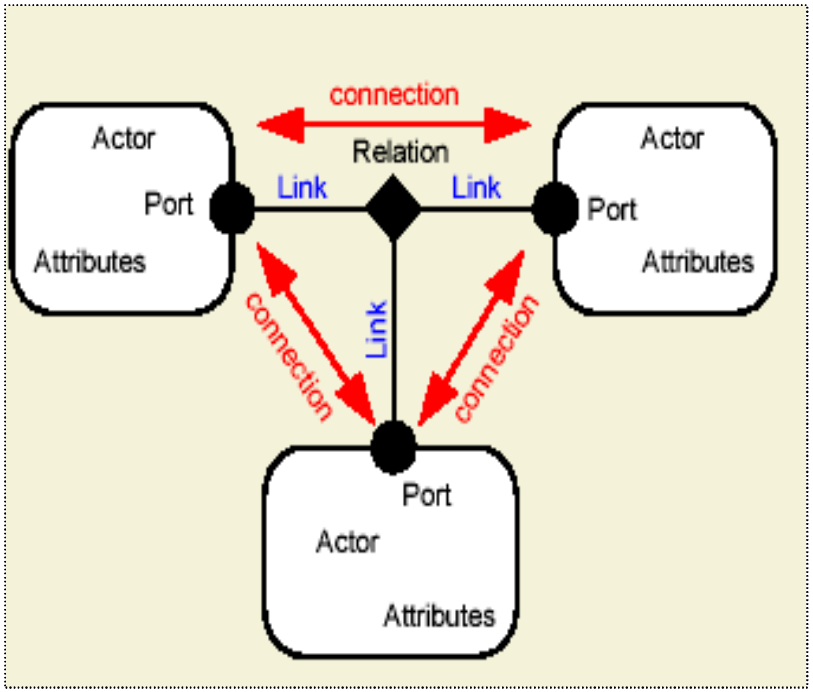

Actor-Oriented Design

- **Actor**
  - Encapsulation of parameterized actions
  - Interface defined by ports and parameters
- **Port**
  - Communication between input and output data
  - Without call-return semantics
- **Relation**
  - Links from output Ports to input Ports
  - Could be 1:1, m:n.
- **Actor Examples**
  - Web service Actor
  - Matlab Actor
  - File Read Actor
  - Local Execution Actor
  - Job Submission Actor
  - …

Adapted from one ppt slides by Edward A. Lee, UC Berkeley

Directors are the Workflow Engines that…

- **Implement different computational models**
- **Define the semantics of**
  - execution of actors and workflows
  - interactions between actors

Ptolemy and Kepler are **unique** in combining multiple and different execution models in one workflow.

• Synchronous Dataflow	• Process Networks
• Dynamic Dataflow	• Rendezvous
• Time Triggered	• Publish and Subscribe
• Synchronous/reactive model	• Continuous Time
• Discrete Event	• Finite State Machines

Kepler Modeling with Graphic User Interface

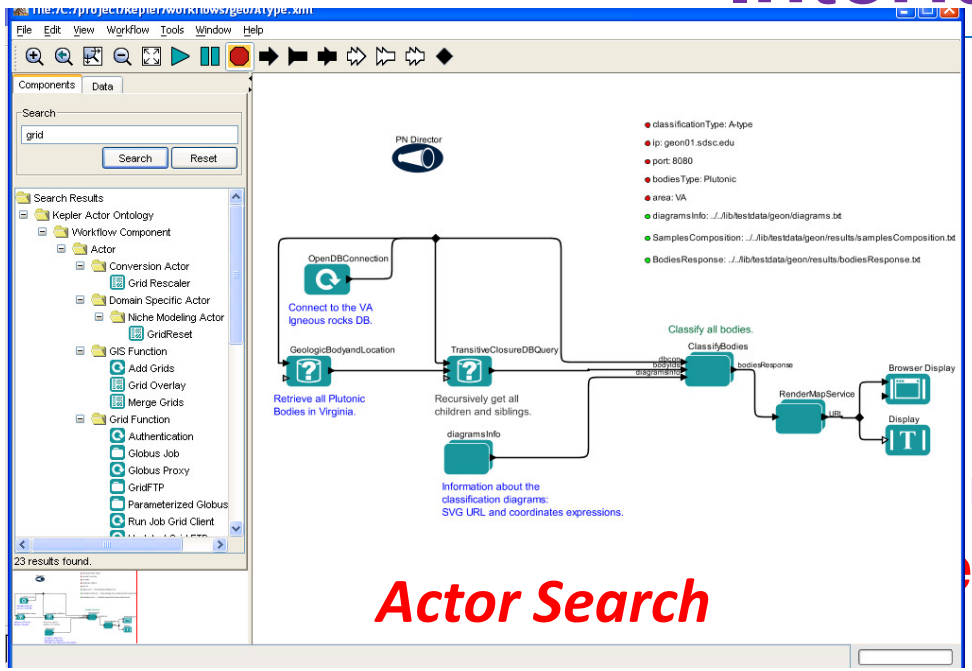

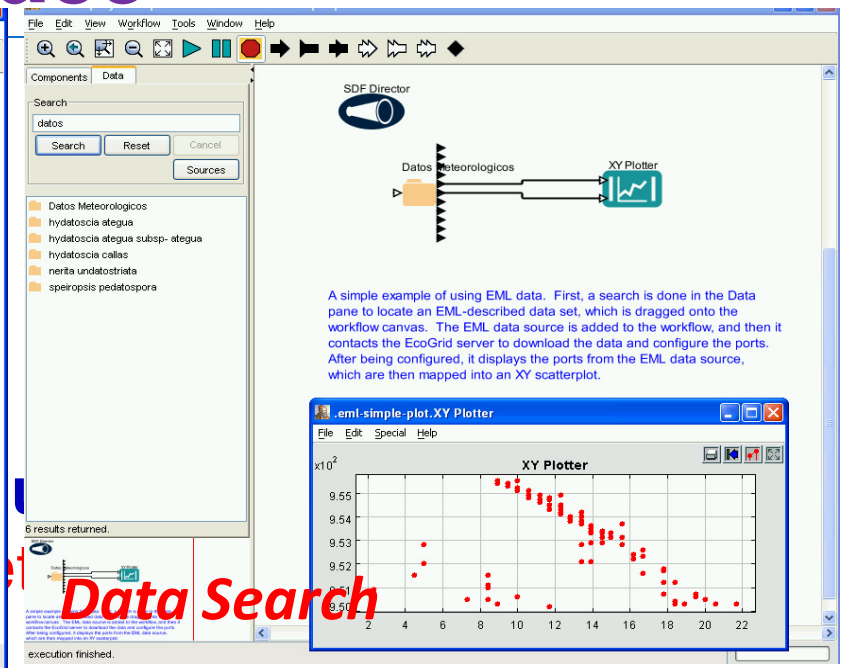

- Actor ontology and semantic search for actors
- Search -> Drag and drop -> Link via ports
- Metadata-based search for datasets

Distributed Execution Requirements and Supportings in Kepler

Requirements	Supportings in Kepler
Ad-hoc network resources	DistributedCompositeActor
Web service resources	Web Service Actors
Cluster resources	Cluster Job Actors
Grid resources	Globus Actors
MapReduce programming model	MapReduce Actor
Cloud resources	Amazon EC2 Actors
…	…

Three Distributed Execution Levels in Kepler

- **Workflow level:** the whole workflow can be executed in distributed environments
  - Example: Web service for Kepler workflow execution
- **Actor level:** distributed computing and data resources can be utilized in actors
  - Example: Cluster Job and Globus actors
- **Sub-workflow level:** sub-workflows can be executed in distributed environments
  - Example: Master-Slave and MapReduce Distributed Execution

Part II

- Kepler Scientific Workflow System
- **Data-Parallel Scientific Workflow Scheduling via Kepler on Virtual Compute Cluster Resources**

Data-Parallel Workflow Scheduling Challenge

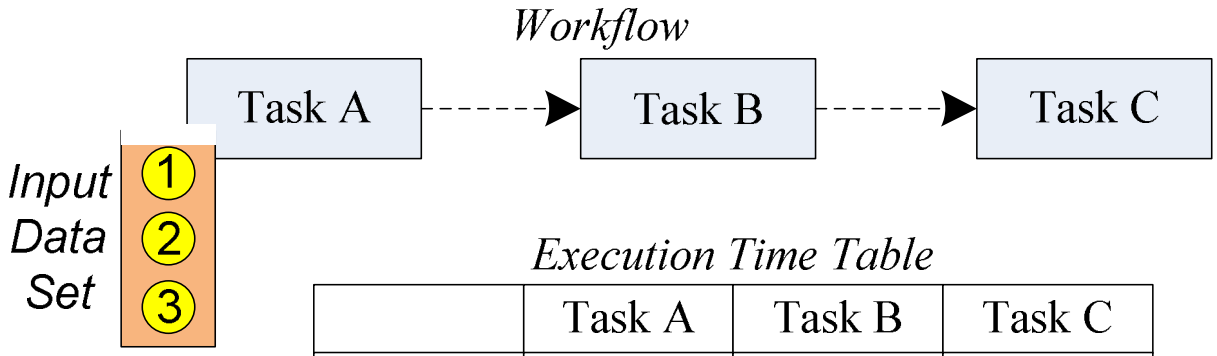

Execution Time Table

	Task A	Task B	Task C
Input 1	10	10	7
Input 2	2	2	2
Input 3	9	6	7

- Data parallel workflow: parallel execution of workflow on multiple parts of input data if these inputs can be processed independently.
- Scheduling challenge on a compute cluster: workflow execution times with different input parts vary and are not predictable.

Data-Parallel Workflow Scheduling Approach in Kepler

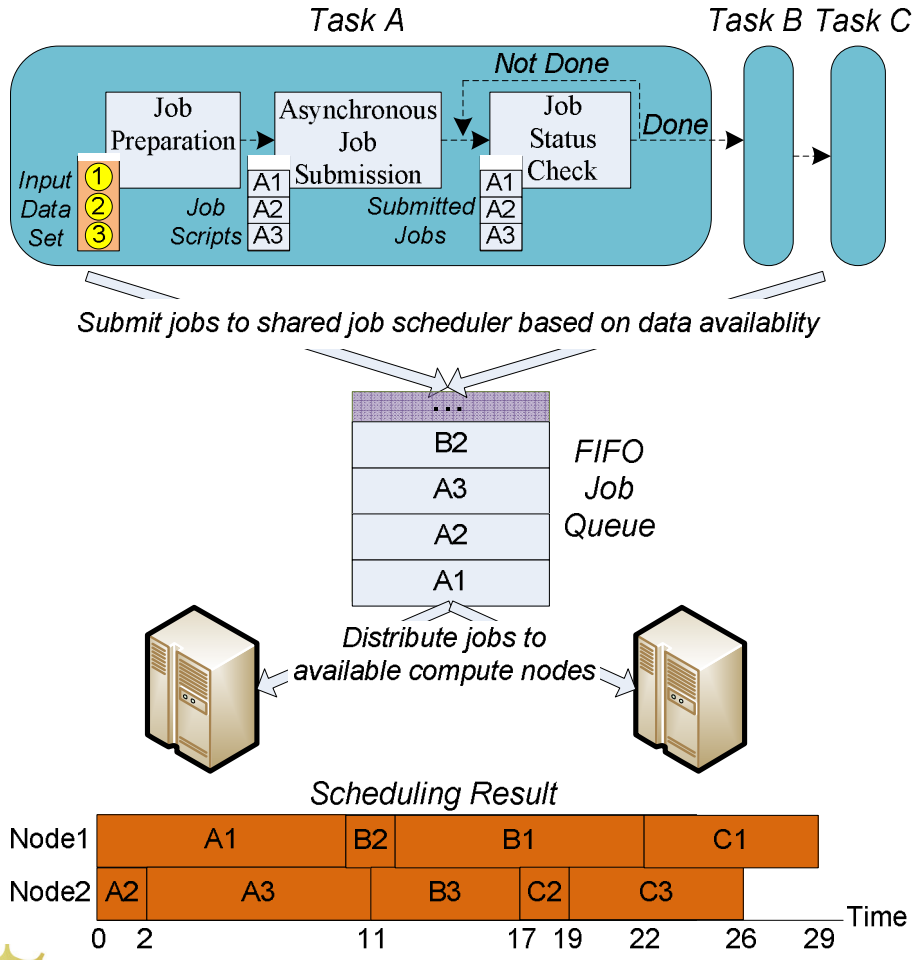

- Data-driven workflow execution (A3 is after A2, B2 is the first in B)
- Pipeline parallelism (B2 and A3)
- Asynchronous job submission
- Shared job scheduler
- Job partition (smaller jobs brings better resource balance)

Application: Enzyme Design Process

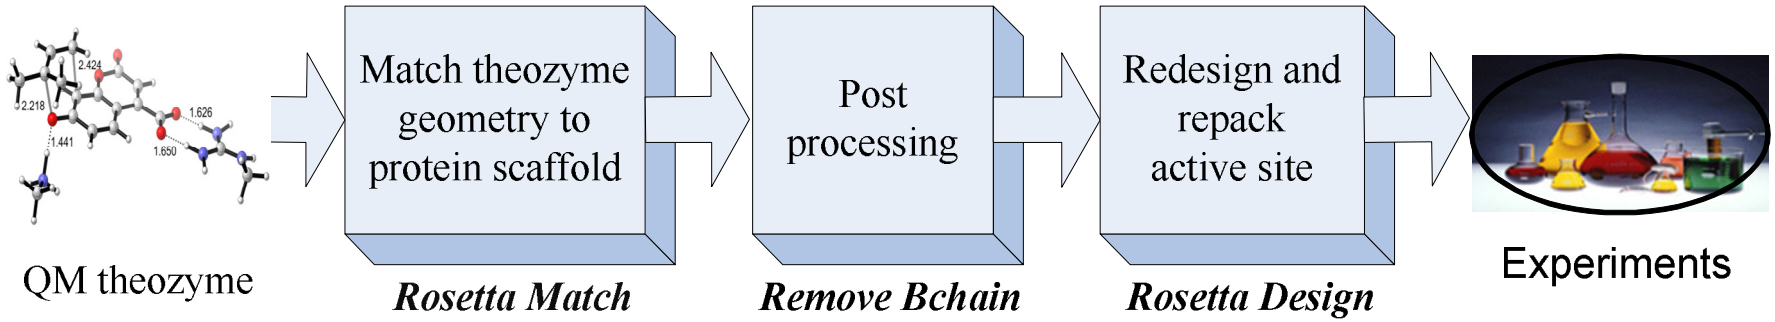

- Data-parallel workflow: 226 inputs totally
- CPU intensive computation
- Computation time for each input varies: seconds to hours

Joint work with Houk's group at UCLA

Kepler Workflow Using Oracle Grid Engine Scheduler

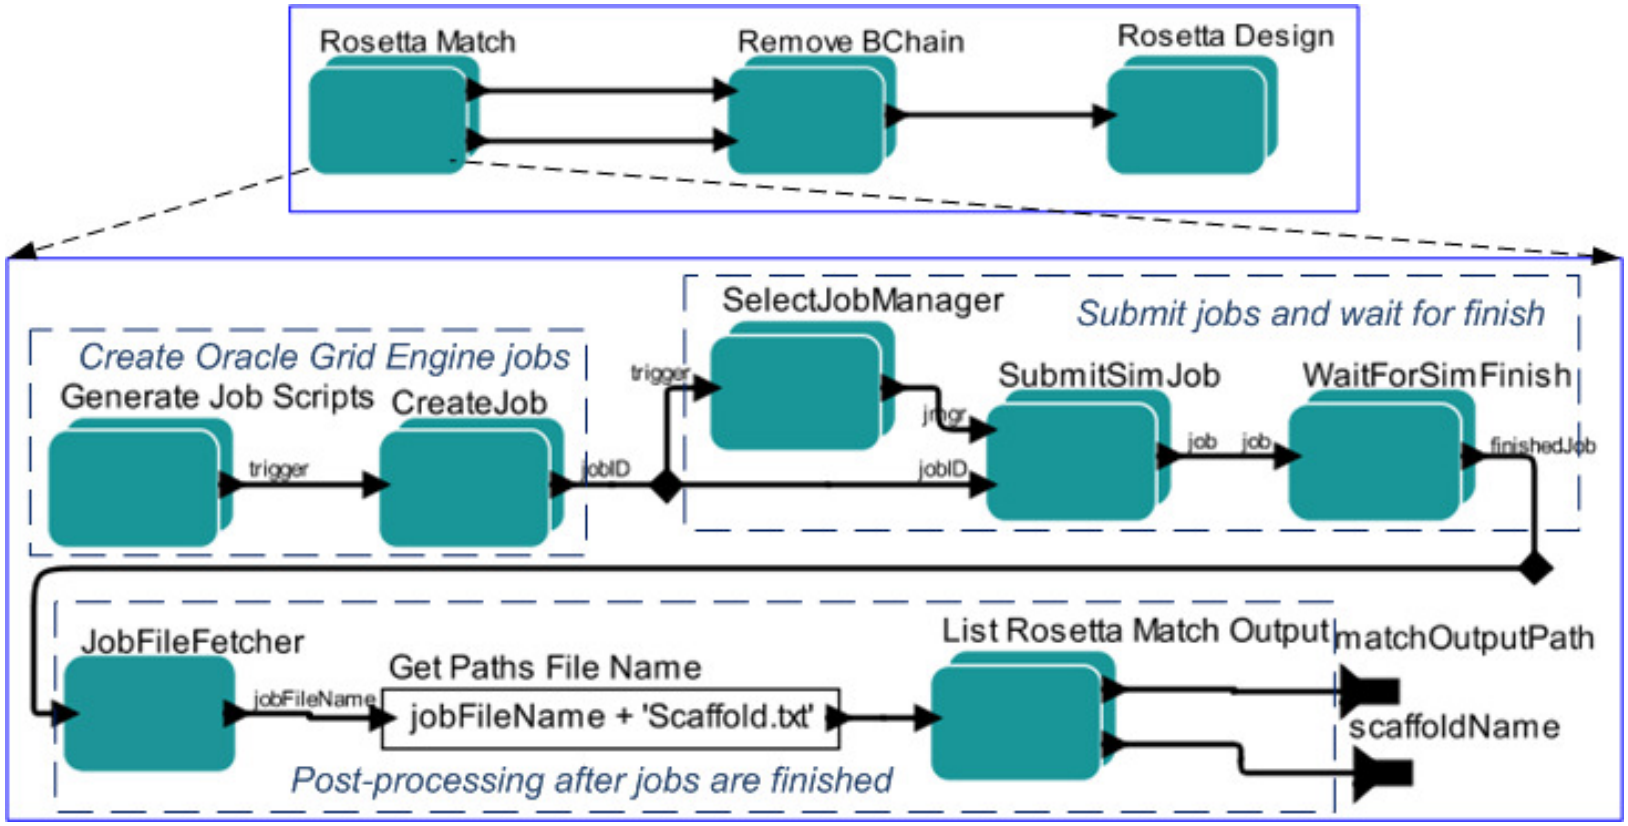

Kepler Workflow Using Hadoop Scheduler

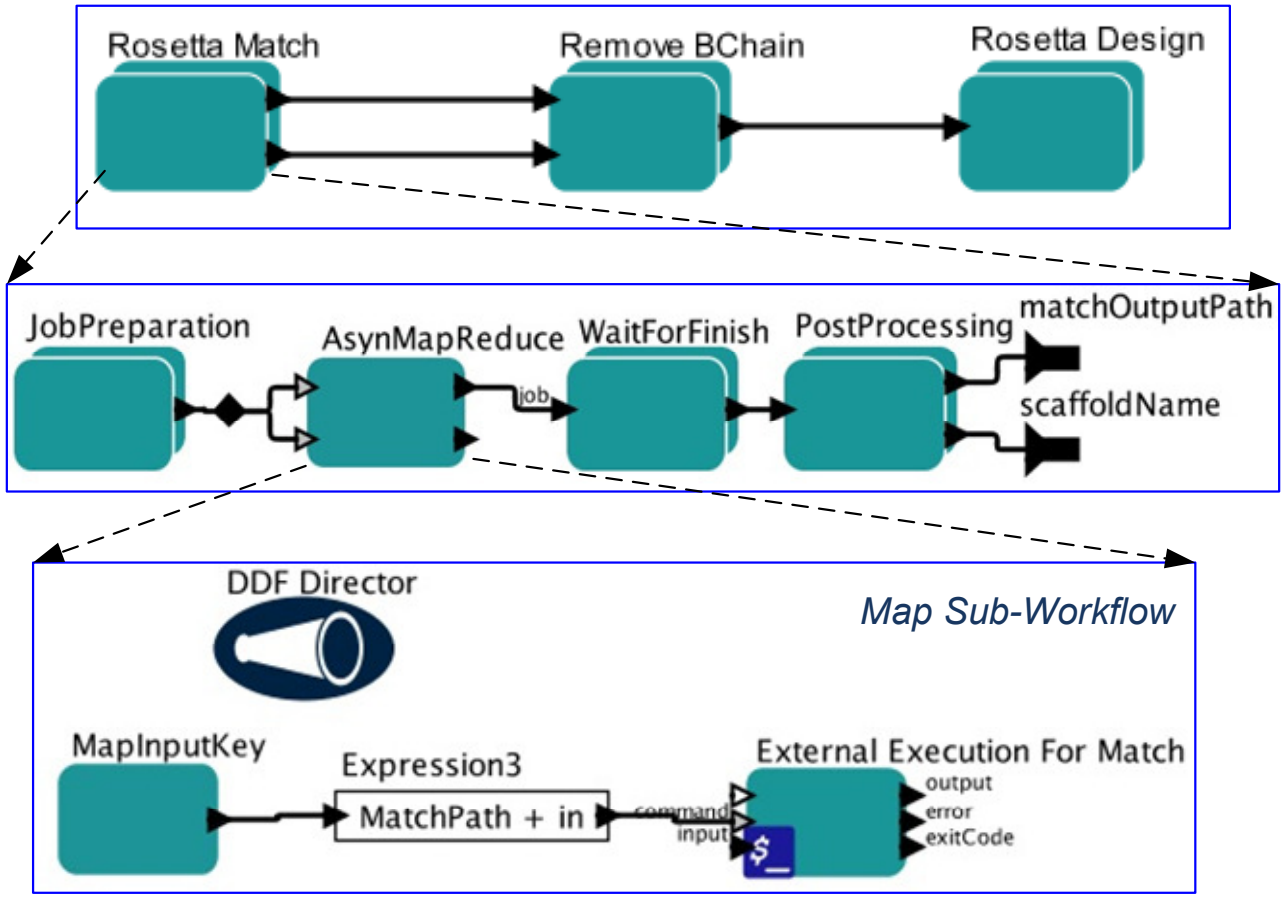

Physical/Virtual Compute Cluster Experimental Environments

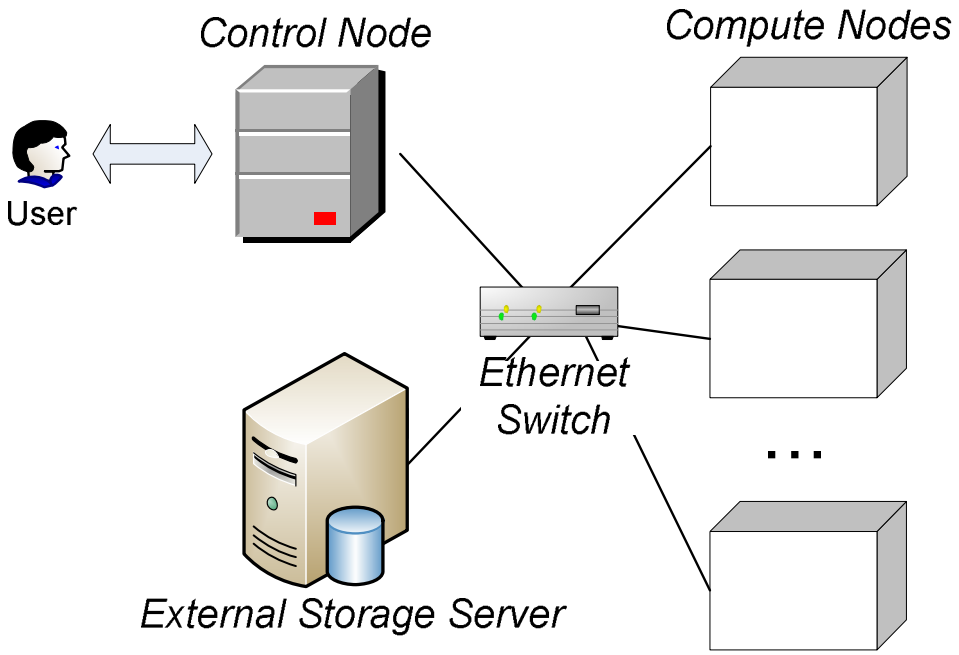

- Physical cluster
- Partially virtual cluster
  - On a physical cluster
  - Only compute nodes virtualized
- Fully virtual cluster
  - On Amazon EC2
  - All components virtualized
  - Setup using StarCluster

Experiments Result

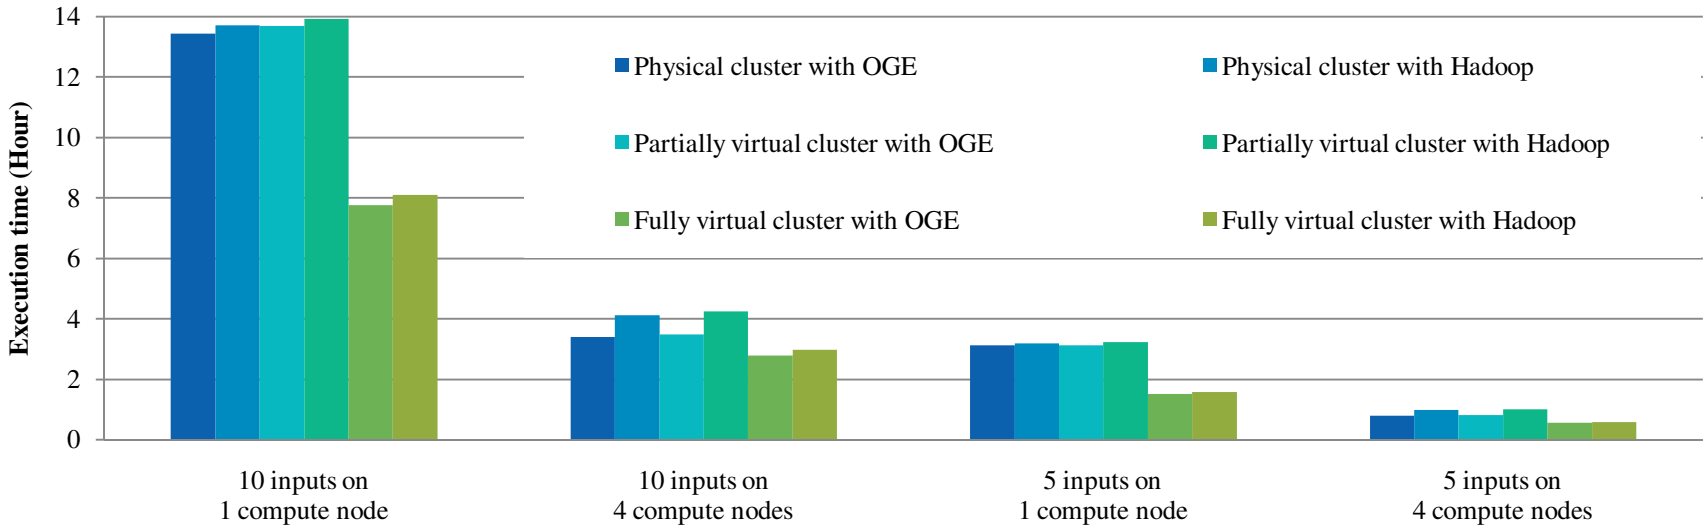

- Execution on the virtual compute cluster only brings very little overhead (around 1%)
- Good execution acceleration rates and resource balance
- The Hadoop workflow execution takes longer time than the Oracle Grid Engine workflow execution
  - Jobs in Hadoop workflow are wrapped by MapReduce sub-workflows
  - Efforts for HDFS and data locality is not utilized in Hadoop Workflow

Thanks!

- **Papers for the Above Work**
  - J. Wang, P. Korambath, I. Altintas. *A Physical and Virtual Compute Cluster Resource Load Balancing Approach to Data-Parallel Scientific Workflow Scheduling.* Accepted by IEEE 2011 Fifth International Workshop on Scientific Workflows (SWF 2011), at 2011 Congress on Services (Services 2011).
  - J. Wang, D. Crawl, I. Altintas. *Kepler + Hadoop – A General Architecture Facilitating Data-Intensive Applications in Scientific Workflow Systems*. In Proc. of the 4th Workshop on Workflows in Support of Large-Scale Science (WORKS09) at SC2009.
  - J. Wang, P. Korambath, S. Kim, S. Johnson, K. Jin, D. Crawl, I. Altintas, S. Smallen, B. Labate, K. N. Houk. *Theoretical Enzyme Design Using the Kepler Scientific Workflows on the Grid*. In Proc. of the 5th Workshop on Computational Chemistry and Its Applications (5th CCA) at ICCS 2010.
- **More Information:**
  - Distributed Execution Interest Group of Kepler: https://dev.kepler-project.org/developers/interest-groups/distributed
  - Contact: jianwu@sdsc.edu