FutureGrid Services II
Using HPC Systems MapReduce & Eucalyptus on FutureGrid

FutureGrid Tutorial at PPAM 2011

Torun Poland

September 11 2011

Gregory G. Pike,
Andrew Younge,
Geoffrey Fox,
Gregor von Laszewski

A brief overview

- FutureGrid is a testbed
  - Varied resources with varied capabilities
  - Support for grid, cloud, HPC
  - Continually evolving
  - Sometimes breaks in strange and unusual ways
- FutureGrid as an experiment
  - We're learning as well
  - Adapting the environment to meet user needs

Getting Started

- Getting an account
- Logging in
- Setting up your environment
- Writing a job script
- Looking at the job queue
- Why won’t my job run?
- Getting your job to run sooner

http://portal.futuregrid.org/manual
http://portal.futuregrid.org/tutorials

Getting an account

- Upload your ssh key to the portal, if you have not done that when you created the portal account
  - Account -> Portal Account
    - Edit the ssh key
    - Include the public portion of your SSH key!
    - Use a passphrase when generating the key!!!!!
- Request a FutureGrid HPC/Nimbus Account
  - Account -> HPC & Nimbus
- This process may take up to 3 days.
  - If it's been longer than a week, send email
  - We do not do any account management over weekends!

Generating an SSH key pair

- For Mac or Linux users
  - `ssh-keygen -t rsa`
  - Copy ~/.ssh/id_rsa.pub to the web form

- For Windows users, this is more difficult
  - Download putty.exe and puttygen.exe
  - Puttygen is used to generate an SSH key pair
    - Run puttygen and click “Generate”
  - The public portion of your key is in the box labeled “SSH key for pasting into OpenSSH authorized_keys file”

Logging in

- You must be logging in from a machine that has your SSH key
- Use the following command (on Linux/OSX):
  - `ssh username@india.futuregrid.org`
- Substitute *username* with your FutureGrid account

Now you are logged in. What is next?

Setting up your environment

- Modules is used to manage your $PATH and other environment variables
- A few common module commands
  - `module avail` – lists all available modules
  - `module list` – lists all loaded modules
  - `module load` – adds a module to your environment
  - `module unload` – removes a module from your environment
  - `module clear` –removes all modules from your environment

Writing a job script

- A job script has PBS directives followed by the commands to run your job
- At least specify –l and –q options
- The rest is a normal bash script, add whatever you want!

```
#!/bin/bash
#PBS -N testjob
#PBS -l nodes=1:ppn=8
#PBS –q batch
#PBS –M username@example.com
##PBS –o testjob.out
#PBS -j oe
#
sleep 60
hostname
echo $PBS_NODEFILE
cat $PBS_NODEFILE
sleep 60
```

Writing a job script

- Use the qsub command to submit your job
  - qsub testjob.pbs
- Use the qstat command to check your job

```
> qsub testjob.pbs
25265.i136


> qstat
Job id     Name         User  Time Use S Queue
---------- ------------ ----- -------- - ------
25264.i136 sub27988.sub inca  00:00:00 C batch
25265.i136 testjob      gpike 0        R batch
```

Looking at the job queue

- Both *qstat* and *showq* can be used to show what's running on the system
- The *showq* command gives nicer output
- The *pbsnodes* command will list all nodes and details about each node
- The *checknode* command will give extensive details about a particular node

Run `module load moab` to add commands to path

Why won’t my job run?

Two common reasons:

- The cluster is full and your job is waiting for other jobs to finish
- You asked for something that doesn’t exist
  - More CPUs or nodes than exist
- The job manager is optimistic!
  - If you ask for more resources than we have, the job manager will sometimes hold your job until we buy more hardware

Why won’t my job run?

- Use the checkjob command to see why your job will not run

```
> checkjob 319285


job 319285

Name: testjob
State: Idle
Creds: user:gpike group:users class:batch qos:od
WallTime: 00:00:00 of 4:00:00
SubmitTime: Wed Dec 1 20:01:42
(Time Queued Total: 00:03:47 Eligible: 00:03:26)

Total Requested Tasks: 320

Req[0] TaskCount: 320 Partition: ALL

Partition List: ALL,s82,SHARED,msm
Flags: RESTARTABLE
Attr: checkpoint
StartPriority: 3
NOTE: job cannot run (insufficient available procs: 312 available)
```

Why won’t my job run?

- If you submitted a job that cannot run, use qdel to delete the job, fix your script, and resubmit the job
  - `qdel 319285`
- If you think your job should run, leave it in the queue and send email
- It’s also possible that maintenance is coming up soon

Making your job run sooner

- In general, specify the minimal set of resources you need
  - Use minimum number of nodes
  - Use the job queue with the shortest max walltime
    - `qstat -Q -f`
  - Specify the minimum amount of time you need for the job
    - `qsub -l walltime=hh:mm:ss`

Example with MPI

- Run through a simple example of an MPI job
  - Ring algorithm passes messages along to each process as a chain or string
  - Use Intel compiler and mpi to compile & run
  - Hands on experience with PBS scripts

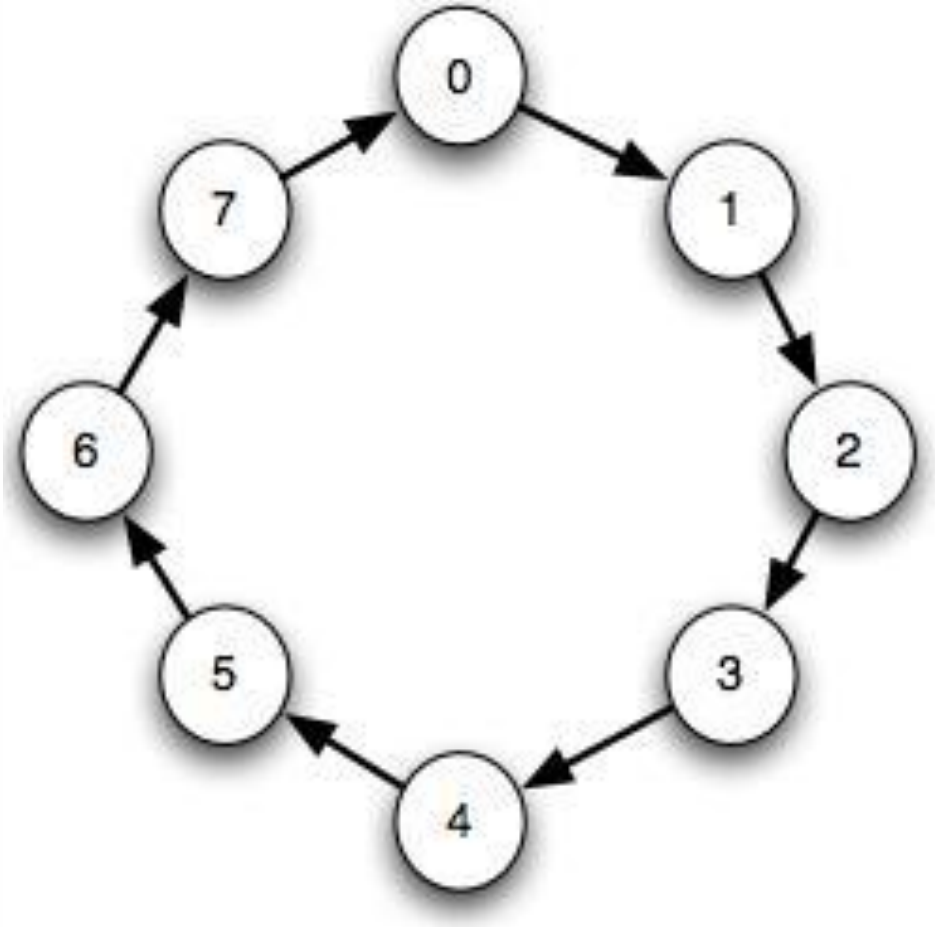

Future Grid

```
#PBS -N hello-mvapich-intel
#PBS -l nodes=4:ppn=8
#PBS -l walltime=00:02:00
#PBS -k oe
#PBS -j oe

EXE=$HOME/mpiring/mpiring

echo "Started on `/bin/hostname`"
echo
echo "PATH is [$PATH]"
echo
echo "Nodes chosen are:"
cat $PBS_NODEFILE
echo
module load intel intelmpi
mpdboot -n 4 -f $PBS_NODEFILE -v --remcons

mpiexec -n 32 $EXE

mpdallexit
```

Lets Run

```
> cp /share/project/mpiexample/mpiring.tar.gz .
> tar xfz mpiring.tar.gz
> cd mpiring
> module load intel intelmpi moab

Intel compiler suite version 11.1/072 loaded
Intel MPI version 4.0.0.028 loaded
moab version 5.4.0 loaded

> mpicc -o mpiring ./mpiring.c
> qsub mpiring.pbs
100506.i136

> cat ~/hello-mvapich-intel.o100506
```

...

MapReduce on FutureGrid

Shava Smallen, Andrew Younge,
Geoffrey Fox

Why MapReduce?

- Largest (in data processed) parallel computing platform today as runs information retrieval engines at Google, Yahoo and Bing.
- Portable to Clouds and HPC systems
- Has been shown to support much data analysis
- It is “disk” (basic MapReduce) or “database” (DrayadLINQ) NOT “memory” oriented like MPI; supports “Data-enabled Science”
- Fault Tolerant and Flexible
- Interesting extensions like Pregel and Twister (Iterative MapReduce)
- Spans Pleasingly Parallel, Simple Analysis (make histograms) to main stream parallel data analysis as in parallel linear algebra
  - Not so good at solving PDE’s

Application Classification: MapReduce and MPI

(a) Map Only	(b) Classic MapReduce	(c) Iterative MapReduce	(d) Loosely Synchronous
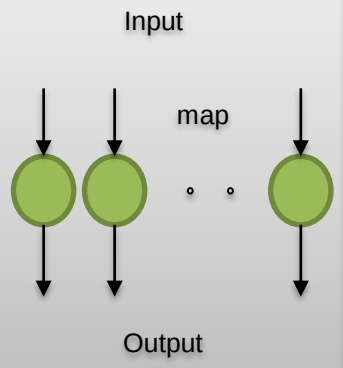	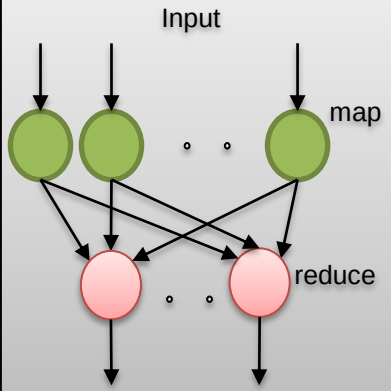	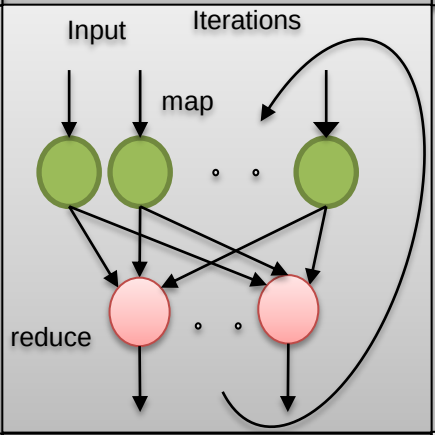	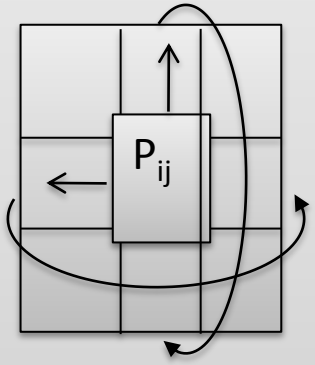
BLAST Analysis Smith-Waterman Distances Parametric sweeps PolarGrid Matlab data analysis	High Energy Physics (HEP) Histograms Distributed search Distributed sorting Information retrieval	Expectation maximization clustering e.g. Kmeans Linear Algebra Multimensional Scaling Page Rank	Many MPI scientific applications such as solving differential equations and particle dynamics
← Domain of MapReduce and Iterative Extensions →			MPI

Microsoft Wants to Make It Easy for Academics to Analyze 'Big Data'

- July 18, 2011, 2:04 pm By Josh Fischman
- http://chronicle.com/blogs/wiredcampus/microsoft-wants-to-make-it-easy-for-academics-to-analyze-big-data/32265
- The enormous amount of data that scholars can generate now can easily overwhelm their desktops and university computing centers. Microsoft Corporation comes riding to the rescue with a new project called Daytona, unveiled at the Microsoft Research Faculty Summit on Monday. Essentially, it's a tool—a free one—that connects these data to Microsoft's giant data centers, and lets scholars run ready-made analytic programs on them. It puts the power of cloud computing at every scholar's fingertips, says Tony Hey, corporate vice president of Microsoft Research Connections, as crunching "Big Data" becomes an essential part of research in health care, education, and the environment.
- Researchers don't need to know how to code for the cloud, for virtual machines, or to write their own software, Mr. Hey says. "What we do needs to be relevant to what academics want," he says, and what they want is to spend time doing research and not writing computer programs. The idea grew out of academe, he adds, with roots in an open-source computing project led by Geoffrey Fox, a professor at Indiana University who directs the Digital Science Center there.
- **This is Iterative MapReduce (aka Twister) on Azure; portably runs on HPC at FutureGrid** **(with Excel front end)**

MapReduce

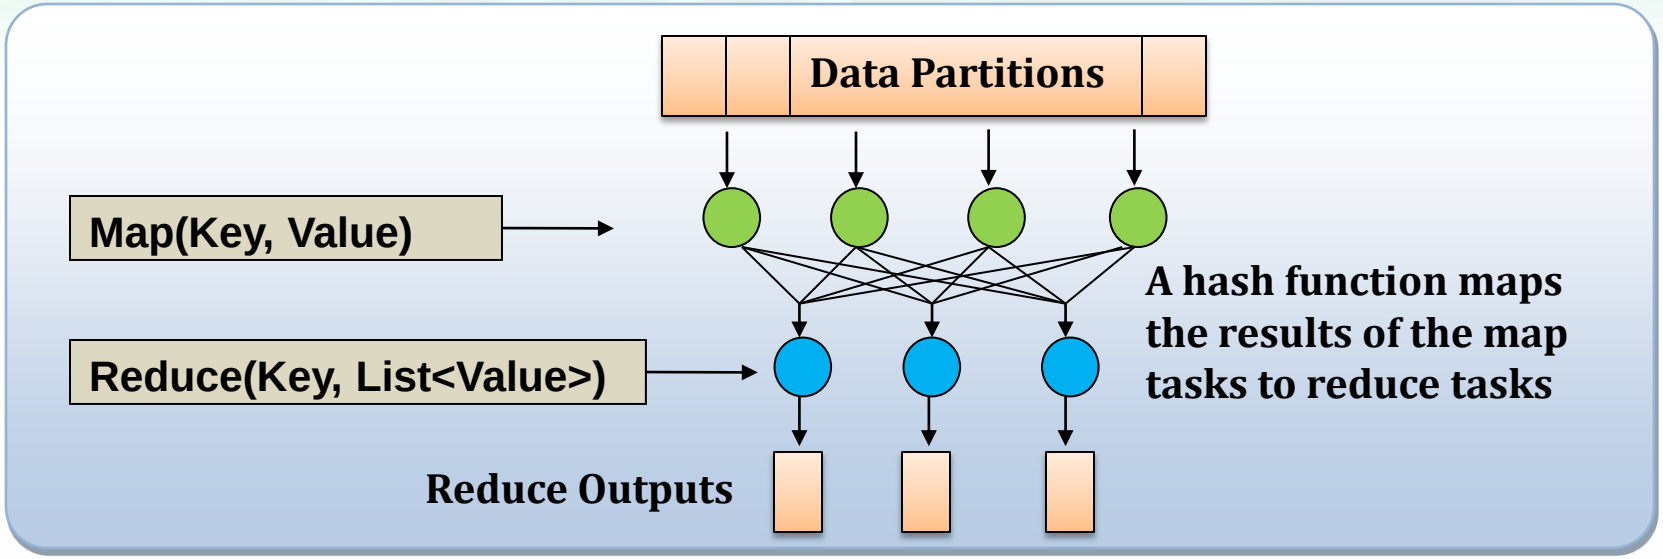

- Implementations (Hadoop – Java; Dryad – Windows) support:
  - Splitting of data with customized file systems
  - Passing the output of map functions to reduce functions
  - Sorting the inputs to the reduce function based on the intermediate keys
  - Quality of service
- 20 petabytes per day (on an average of 400 machines) processed by Google using MapReduce September 2007

MapReduce: The Map Step

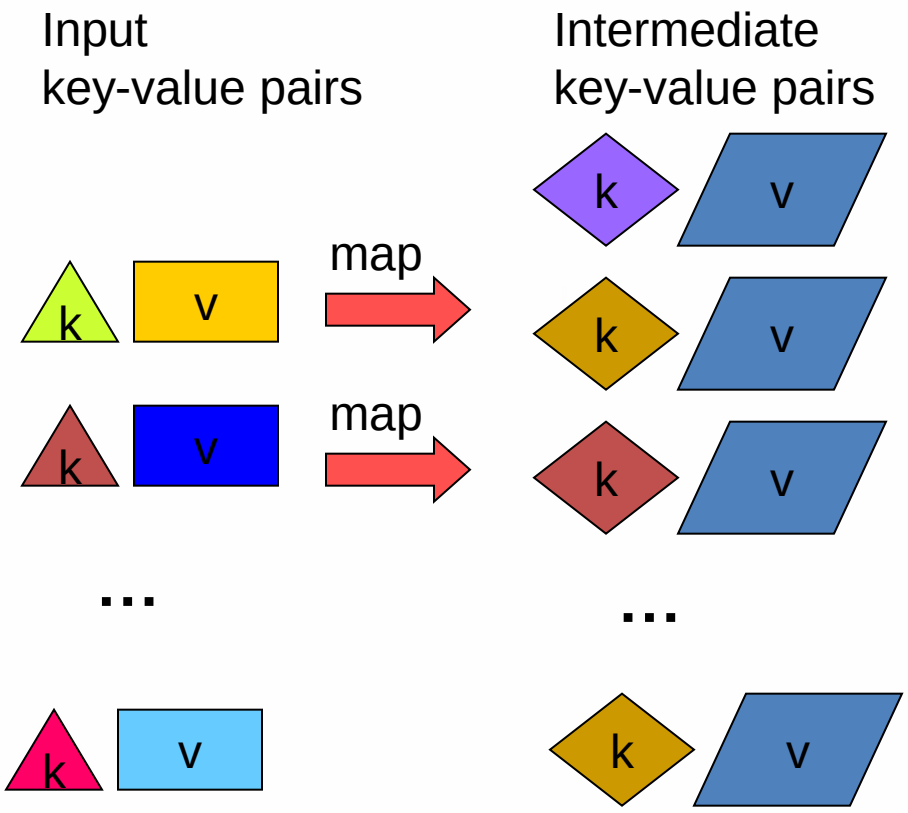

The Map (Example)

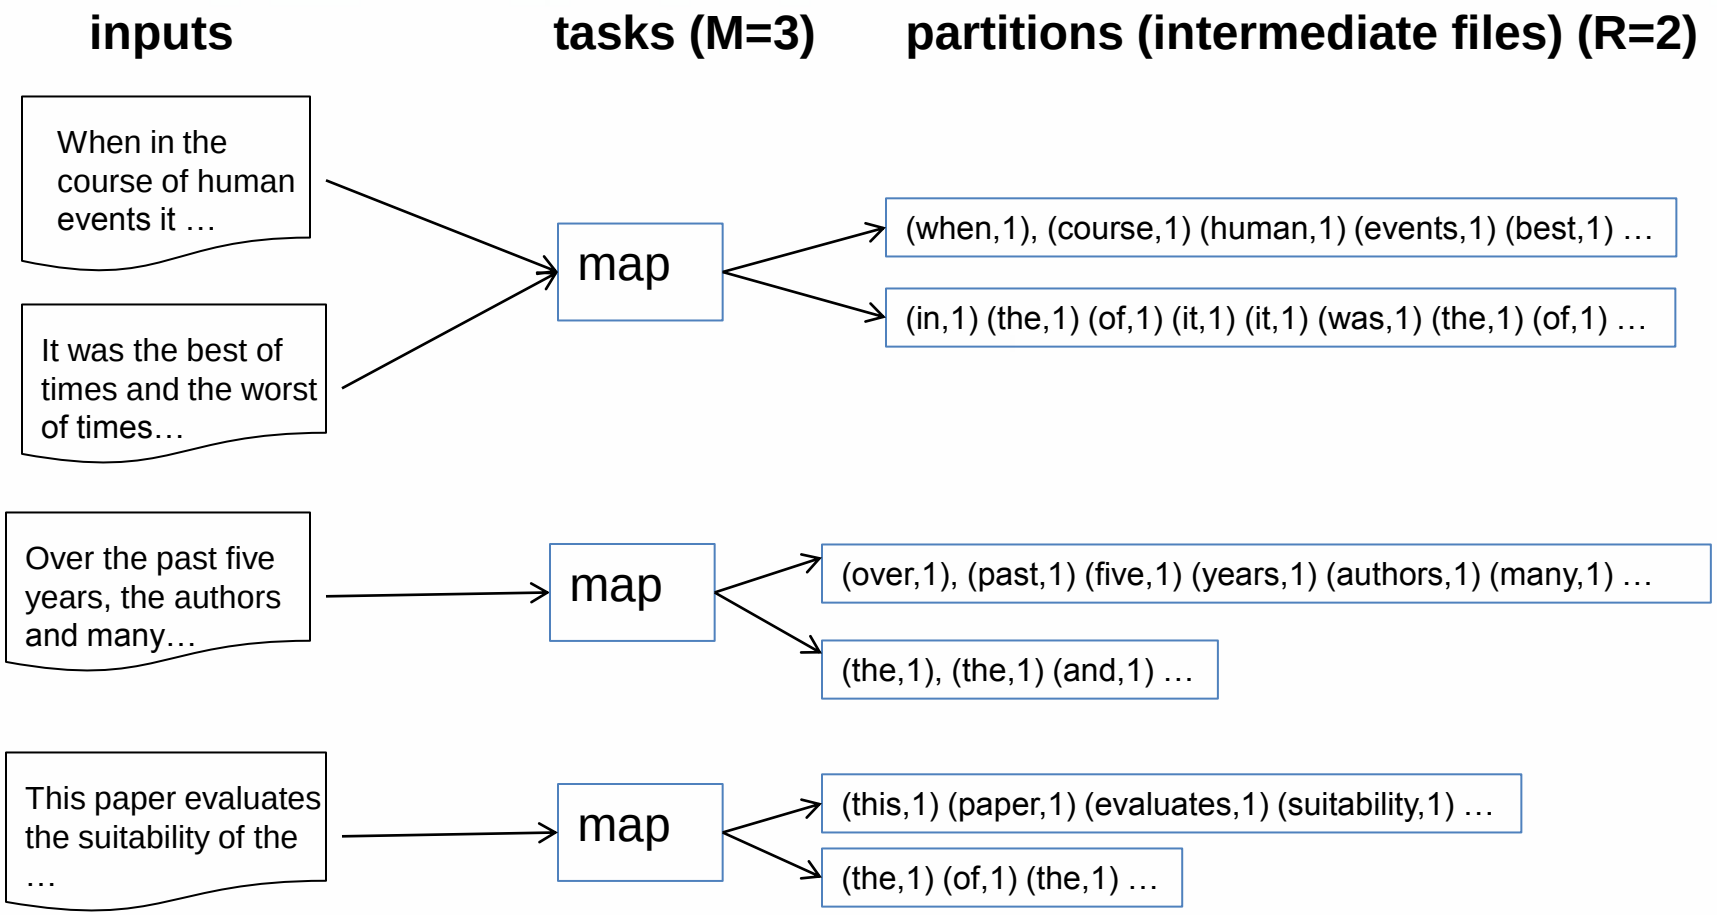

Future Grid

MapReduce: The Reduce Step

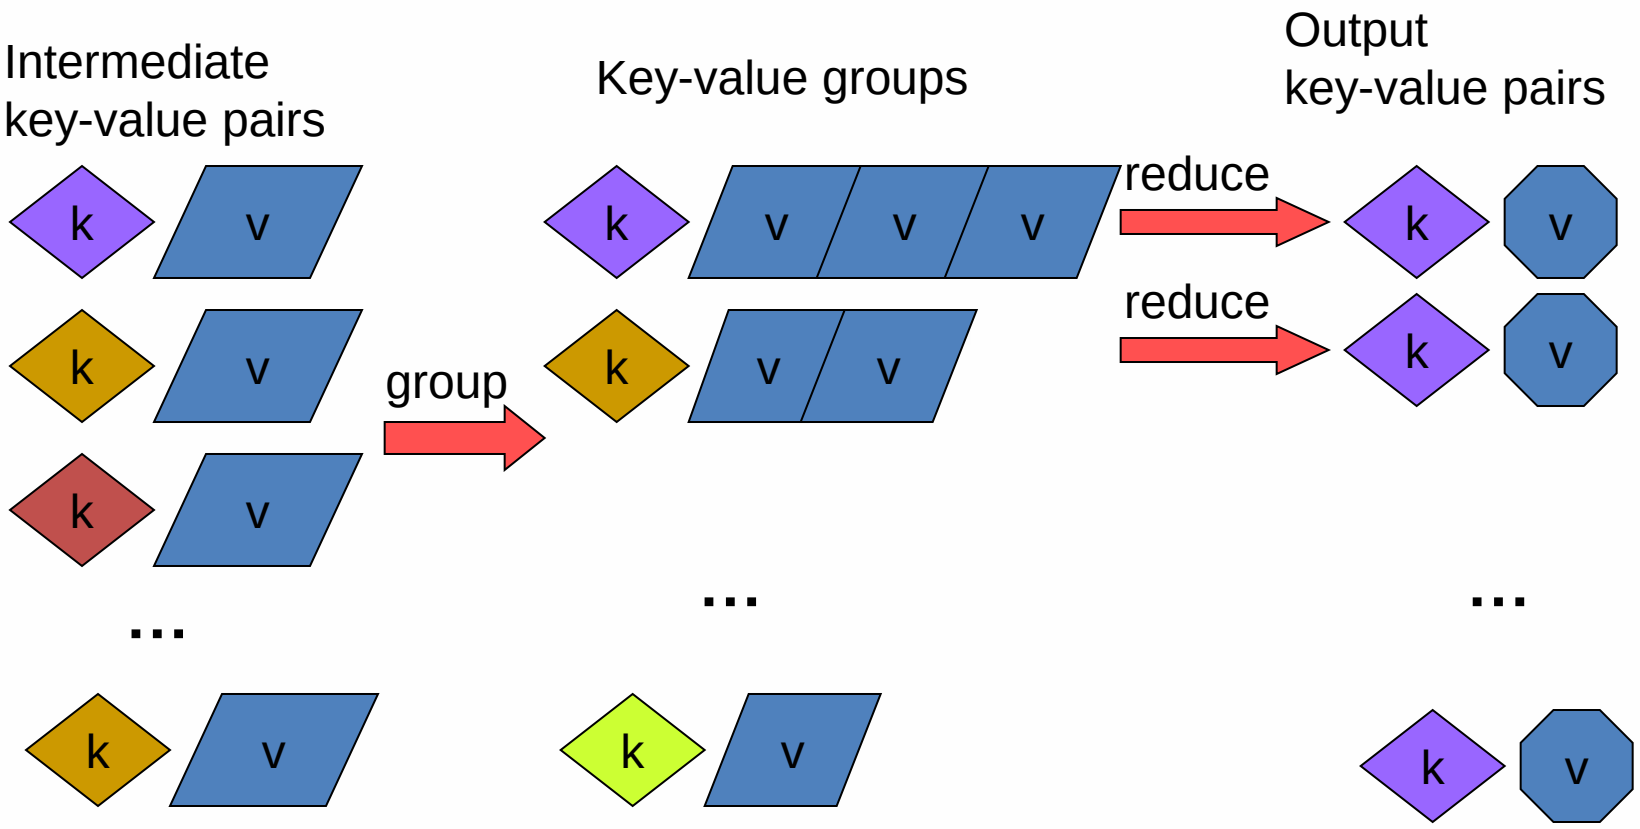

The Reduce (Example)

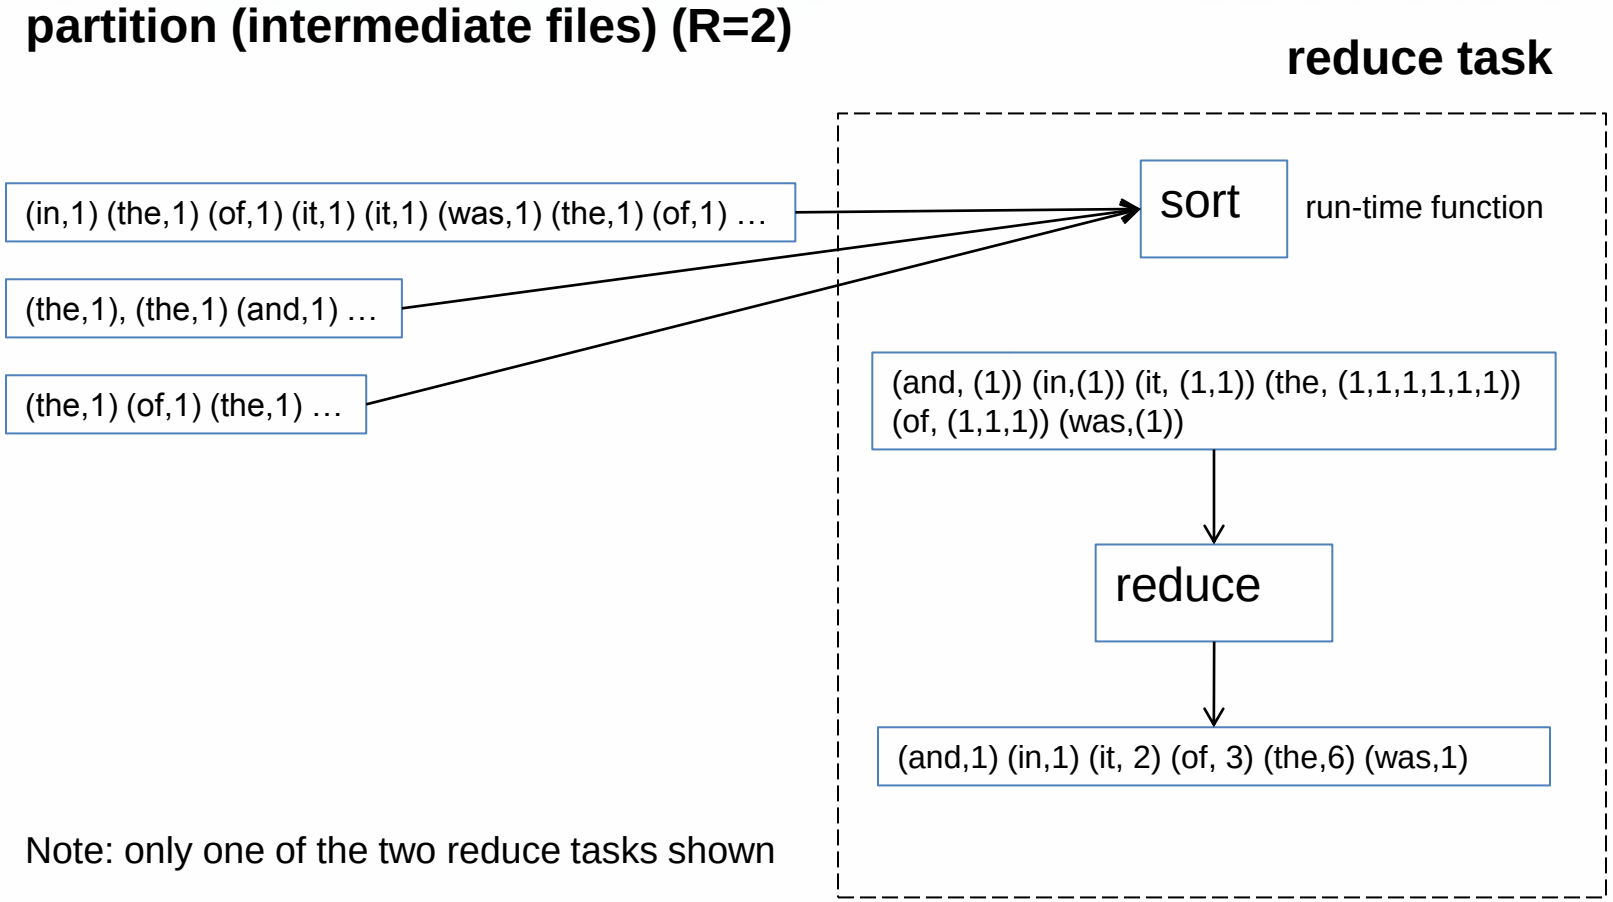

Note: only one of the two reduce tasks shown

Future Grid

Generalizing Information Retrieval

- But you input anything from genome sequences to HEP events as well as documents
- You can map them with an arbitrary program
- You can reduce with an arbitrary reduction including all of those in MPI_(ALL)REDUCE
- In Twister you can iterate this

MapReduce "File/Data Repository" Parallelism

Map = (data parallel) computation reading and writing data

Reduce = Collective/Consolidation phase e.g. forming multiple global sums as in histogram

Instruments

Disks

MPI or Iterative MapReduce

Map Reduce Map Reduce Map

Portals /Users

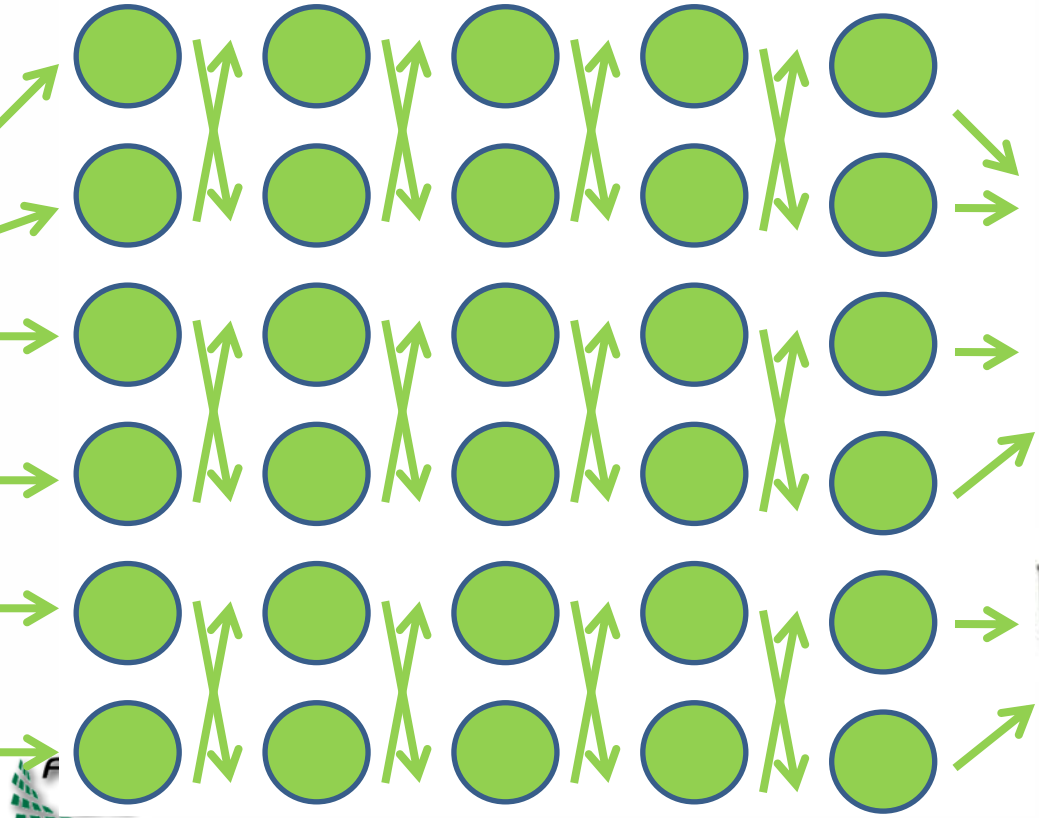

High Energy Physics Data Analysis

An application analyzing data from Large Hadron Collider
(1TB but 100 Petabytes eventually)

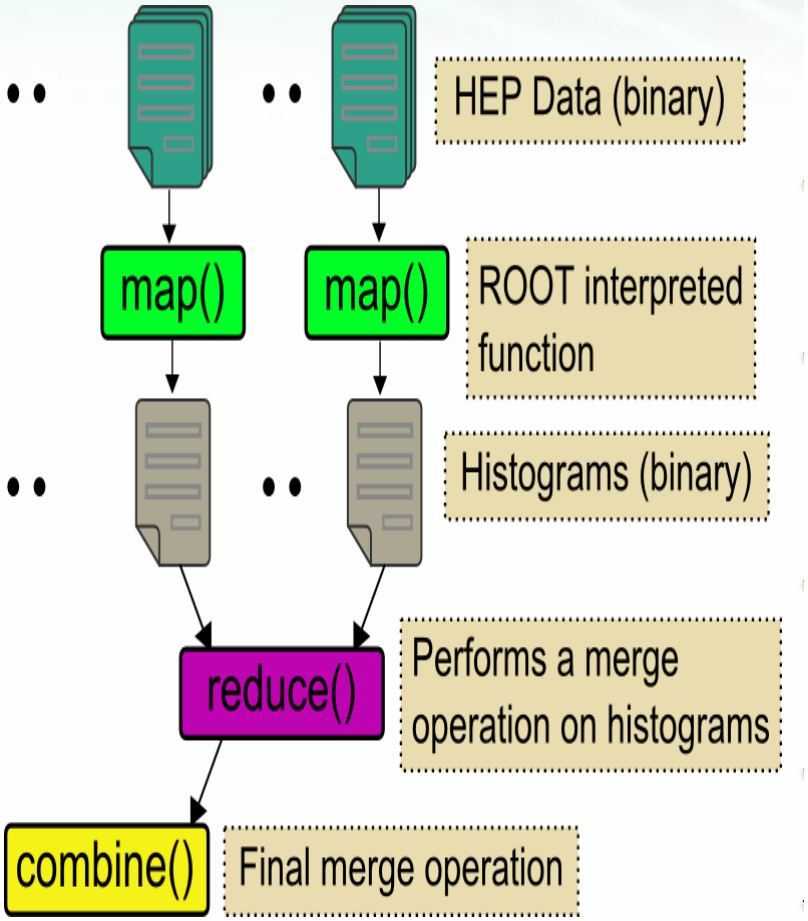

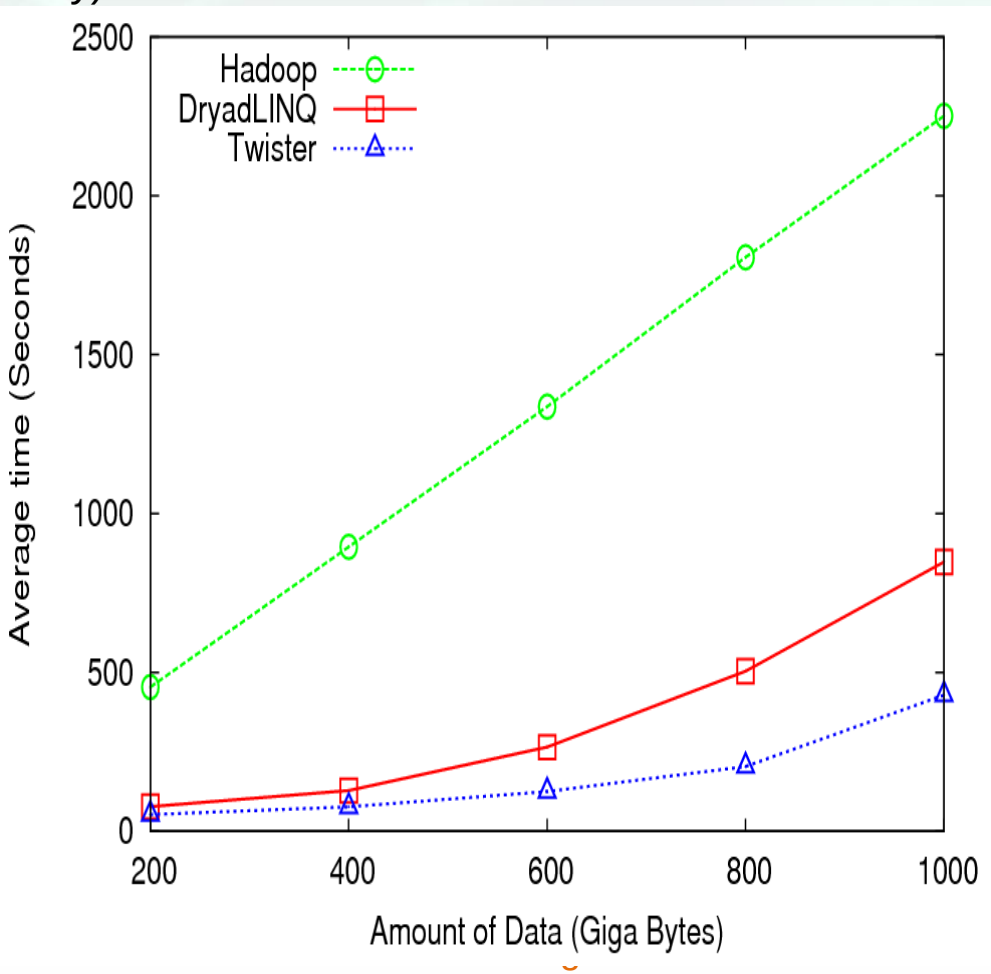

Future Grid

Cap3 Cost

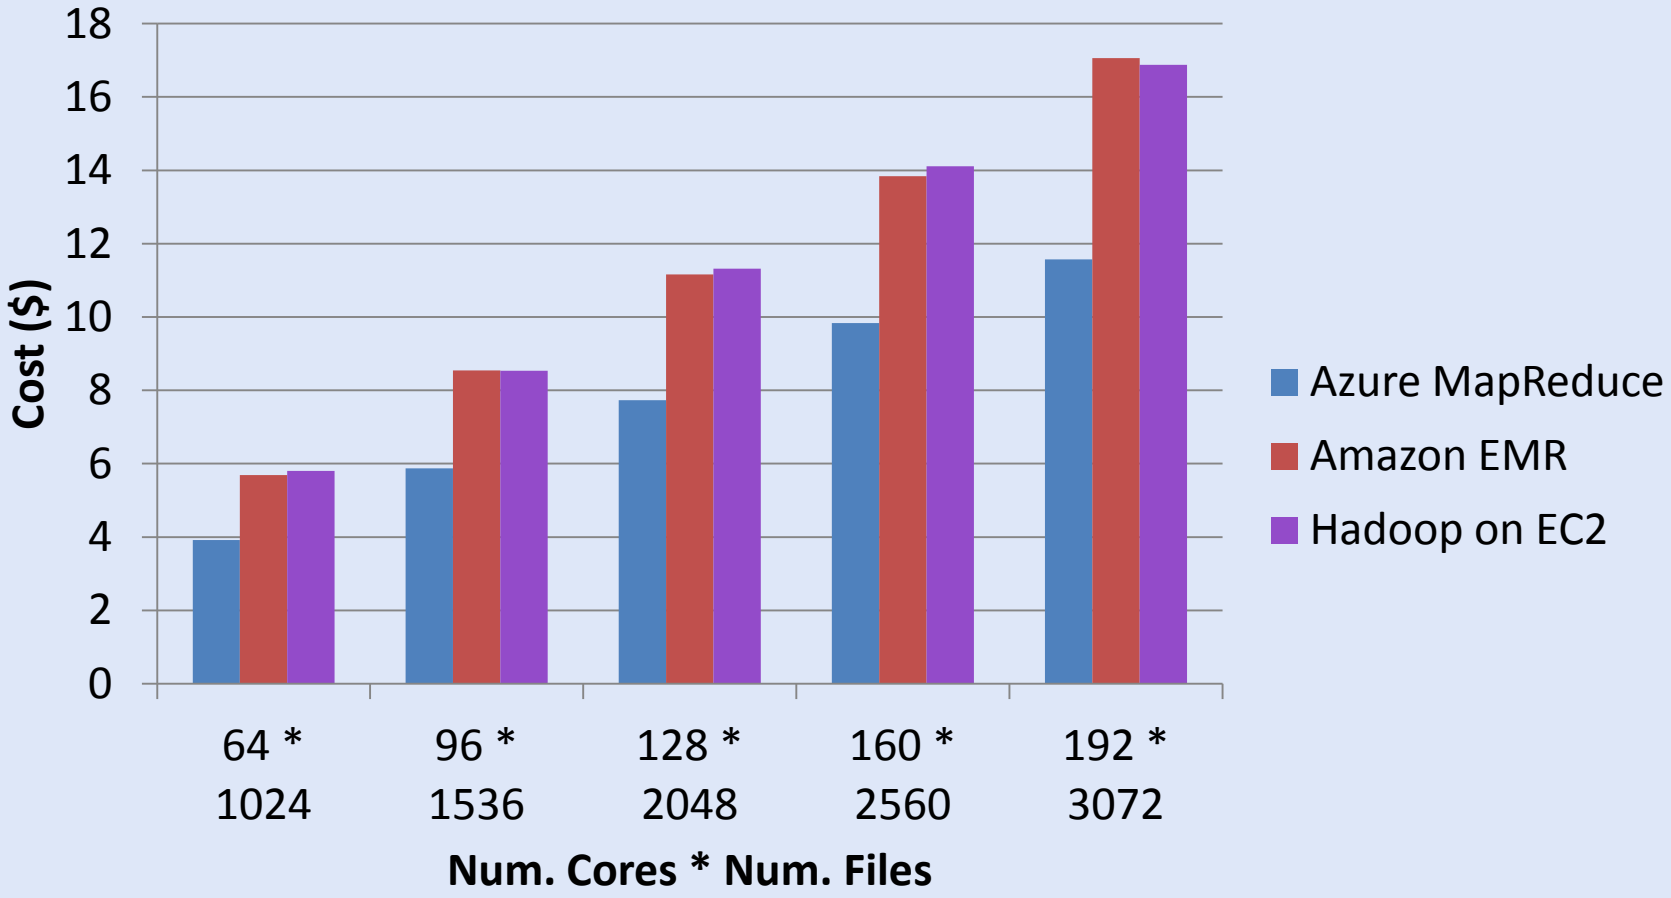
18
16
14
12
10
8
6
4
2
0
Cost ($)
64 * 1024
96 * 1536
128 * 2048
160 * 2560
192 * 3072
Num. Cores * Num. Files
Azure MapReduce
Amazon EMR
Hadoop on EC2

SWG Sequence Alignment Performance

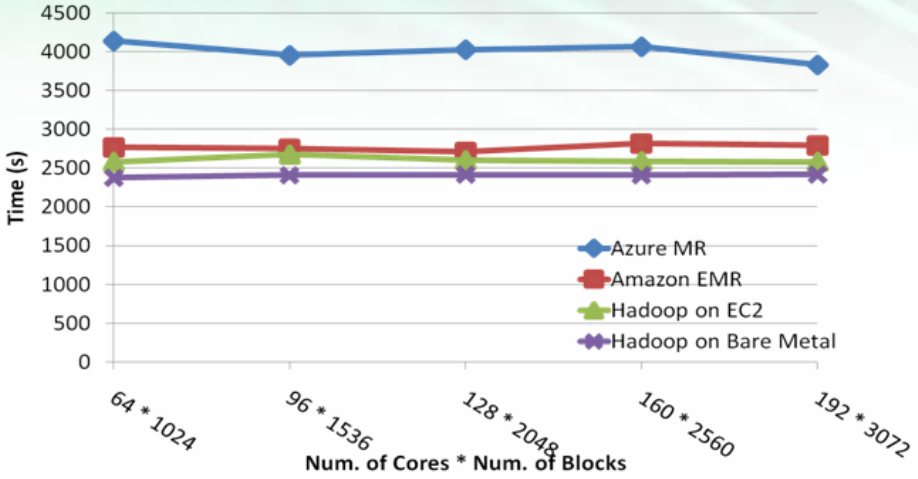

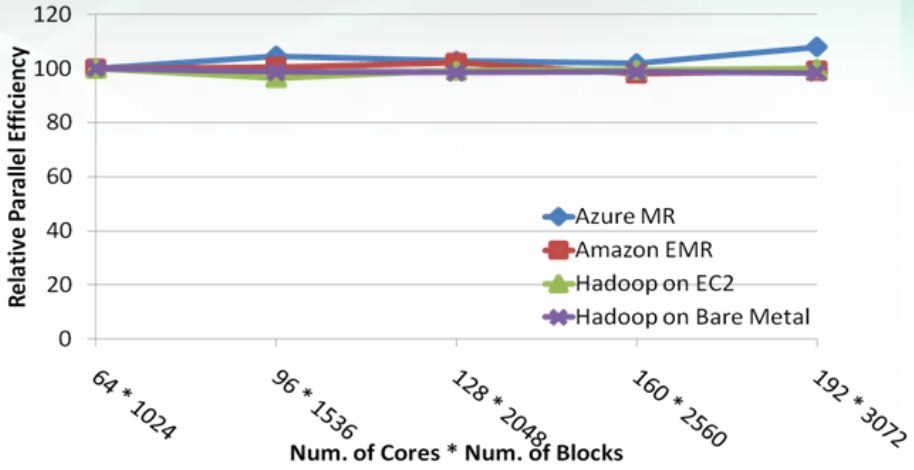

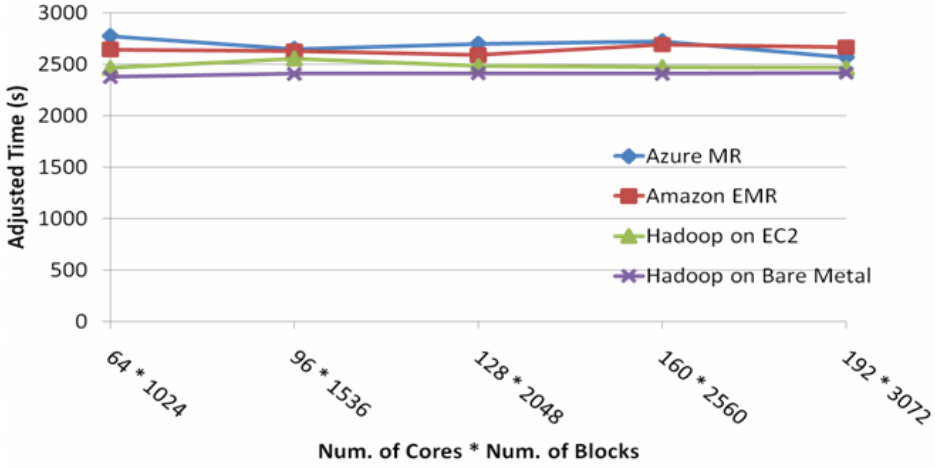

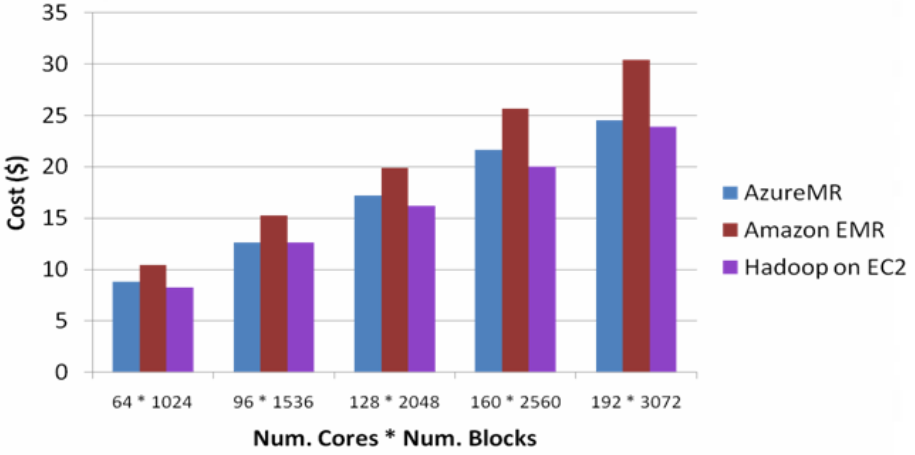

Smith-Waterman-GOTOH to calculate all-pairs dissimilarity

Future Grid

Twister v0.9

March 15, 2011

New Interfaces for Iterative MapReduce Programming
http://www.iterativemapreduce.org/

SALSA Group

Bingjing Zhang, Yang Ruan, Tak-Lon Wu, Judy Qiu, Adam Hughes, Geoffrey Fox, **Applying Twister to Scientific Applications**, Proceedings of IEEE CloudCom 2010 Conference, Indianapolis, November 30-December 3, 2010

Twister4Azure released May 2011
http://salsahpc.indiana.edu/twister4azure/
MapReduceRoles4Azure available for some time at
http://salsahpc.indiana.edu/mapreduceroles4azure/

K-Means Clustering

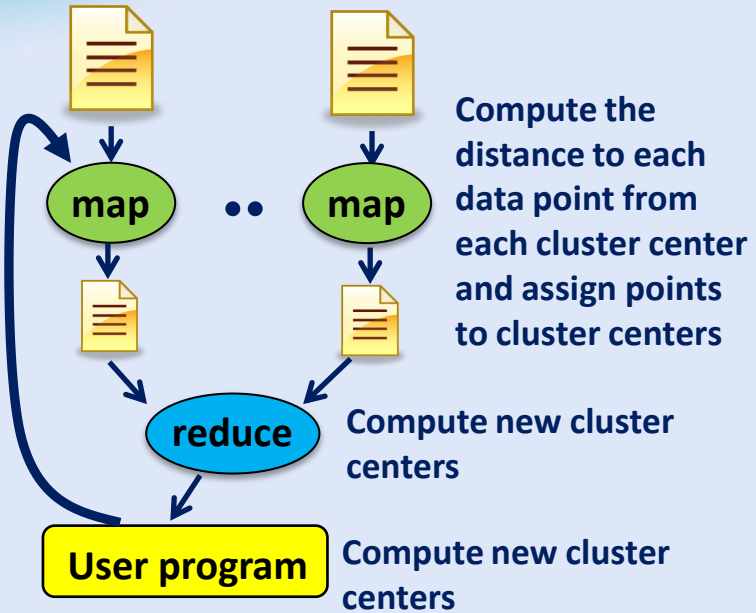

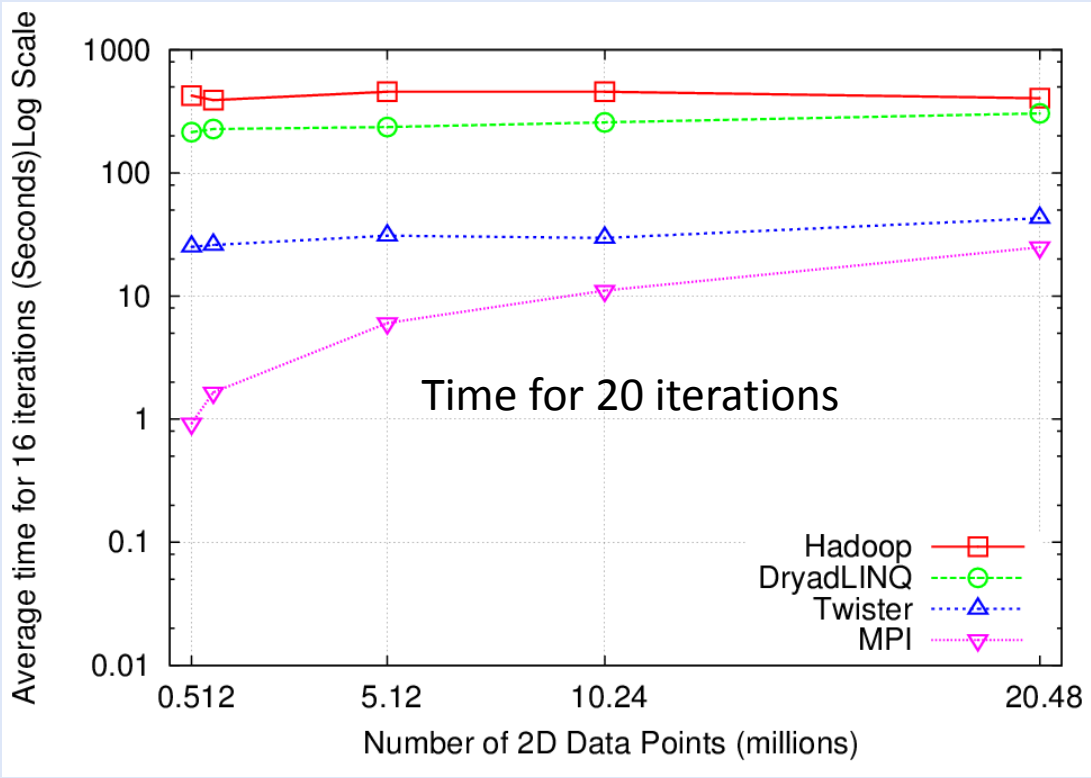

- Iteratively refining operation
- Typical MapReduce runtimes incur extremely high overheads
  - New maps/reducers/vertices in every iteration
  - File system based communication
- Long running tasks and faster communication in Twister enables it to perform close to MPI

Twister

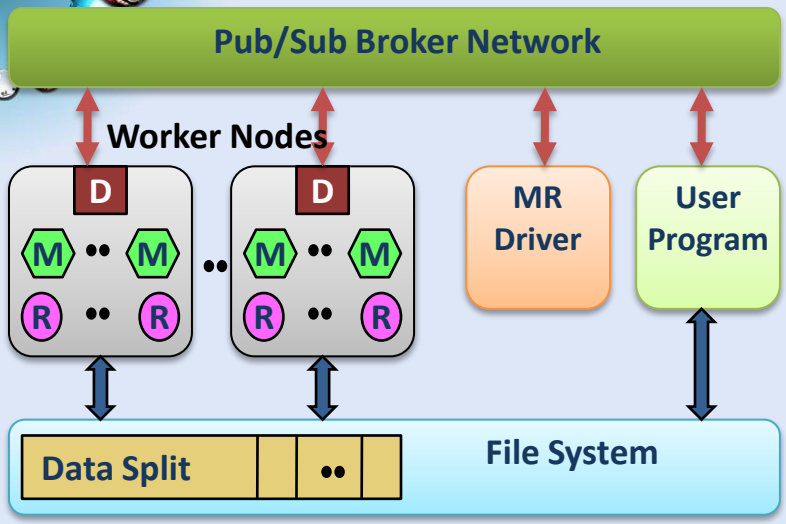

- Streaming based communication
- Intermediate results are directly transferred from the map tasks to the reduce tasks – **eliminates local files**
- Cacheable map/reduce tasks
  - Static data remains in memory
- Combine phase to combine reductions
- User Program is the **composer** of MapReduce computations
- **Extends the MapReduce model to iterative computations**

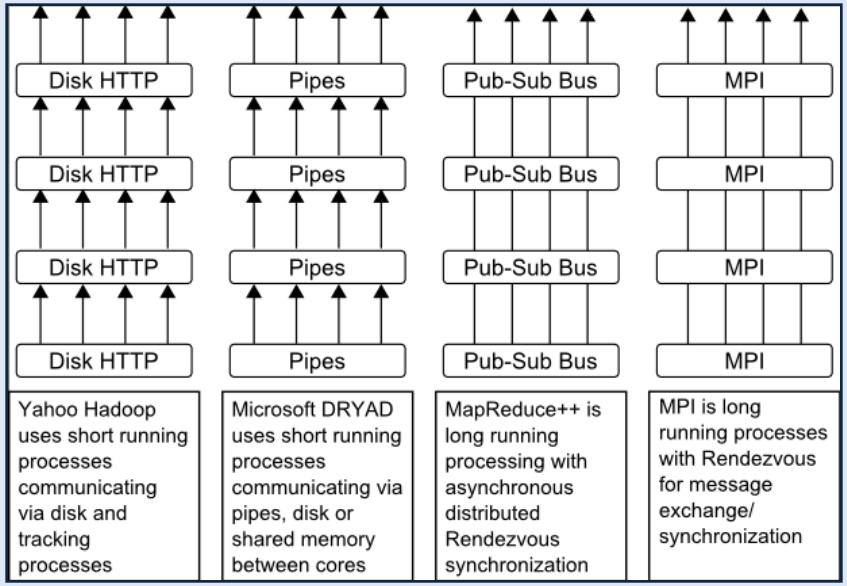

Different synchronization and intercommunication mechanisms used by the parallel runtimes

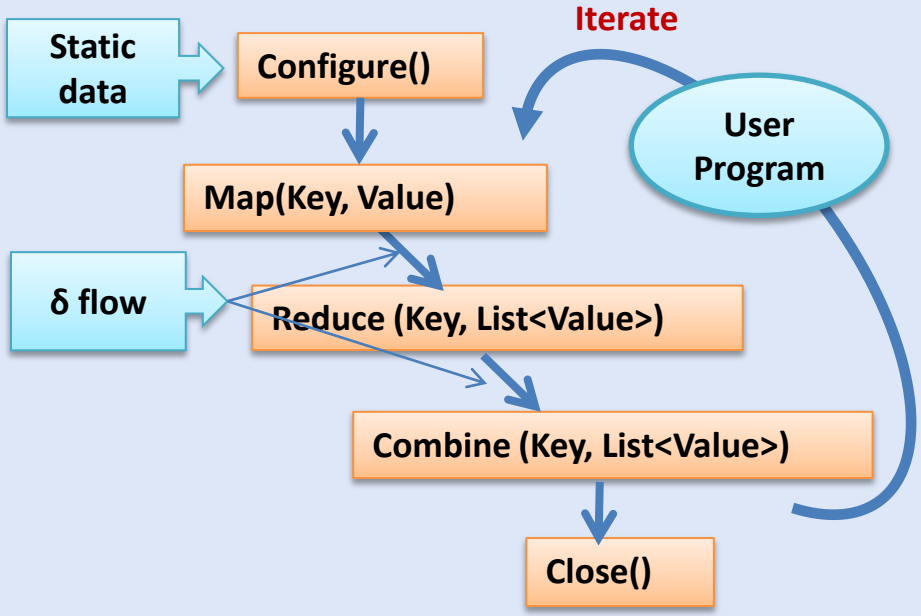

Performance of Pagerank using ClueWeb Data (Time for 20 iterations)

using 32 nodes (256 CPU cores) of Crevasse

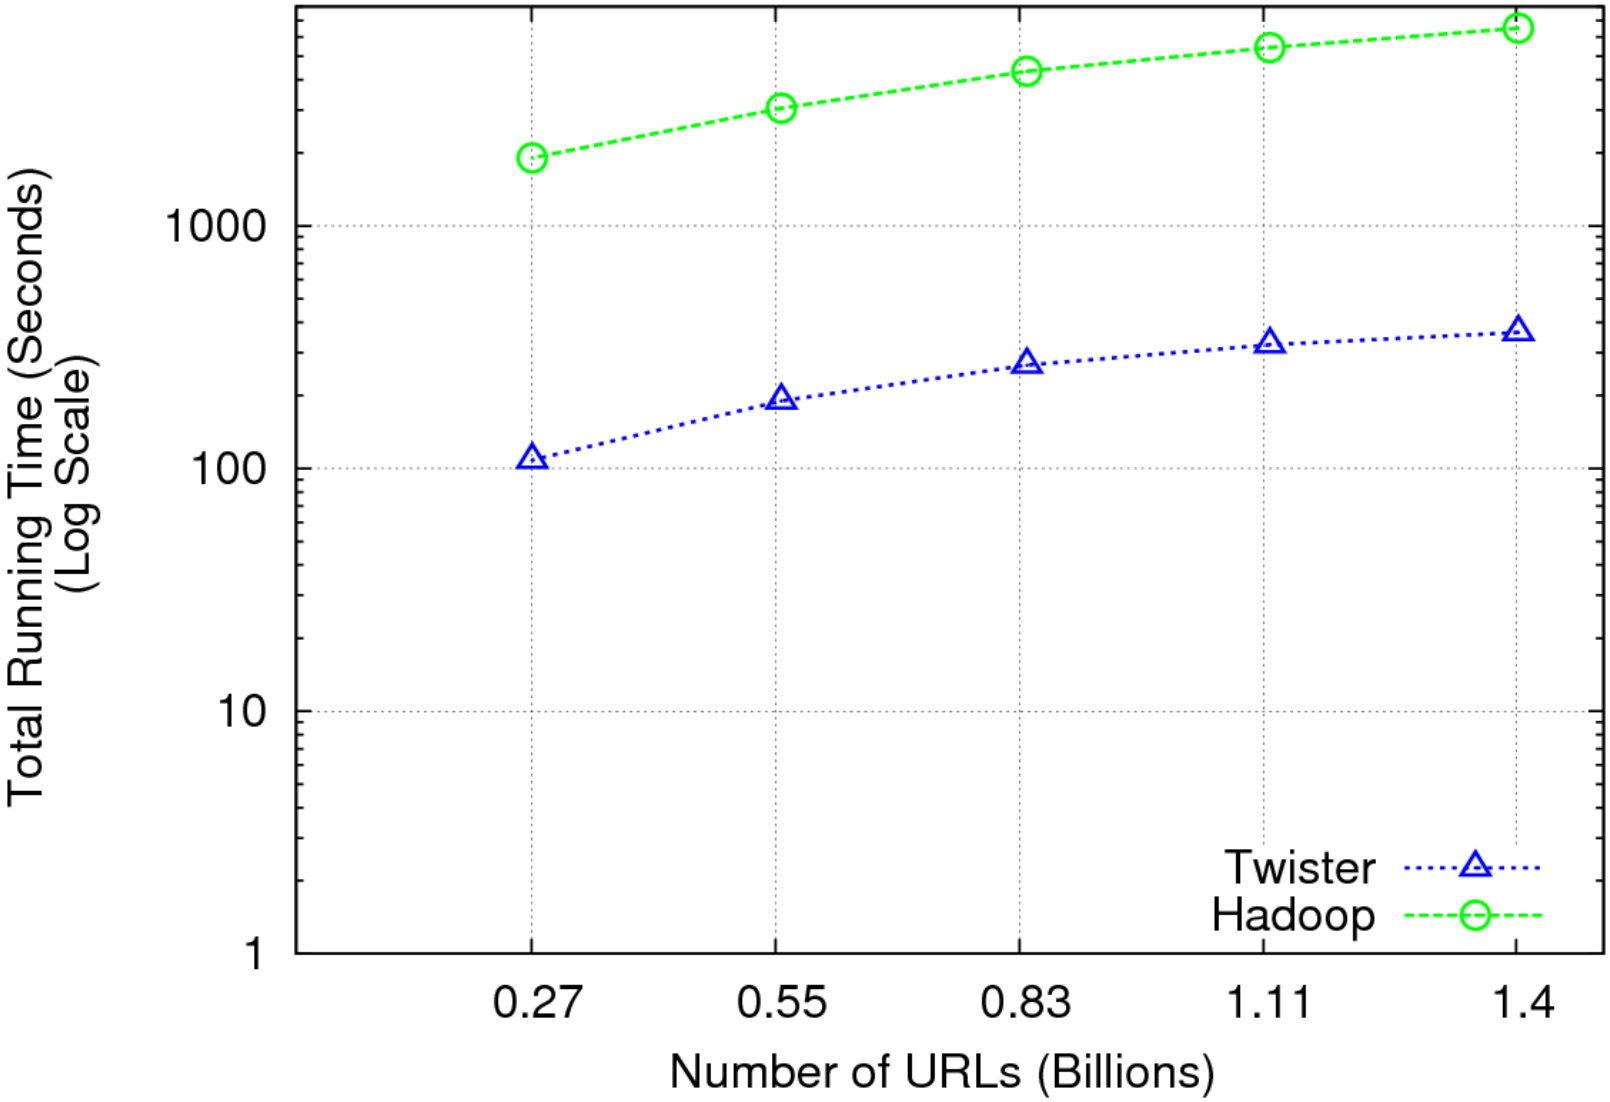

Iterative MapReduce for Azure

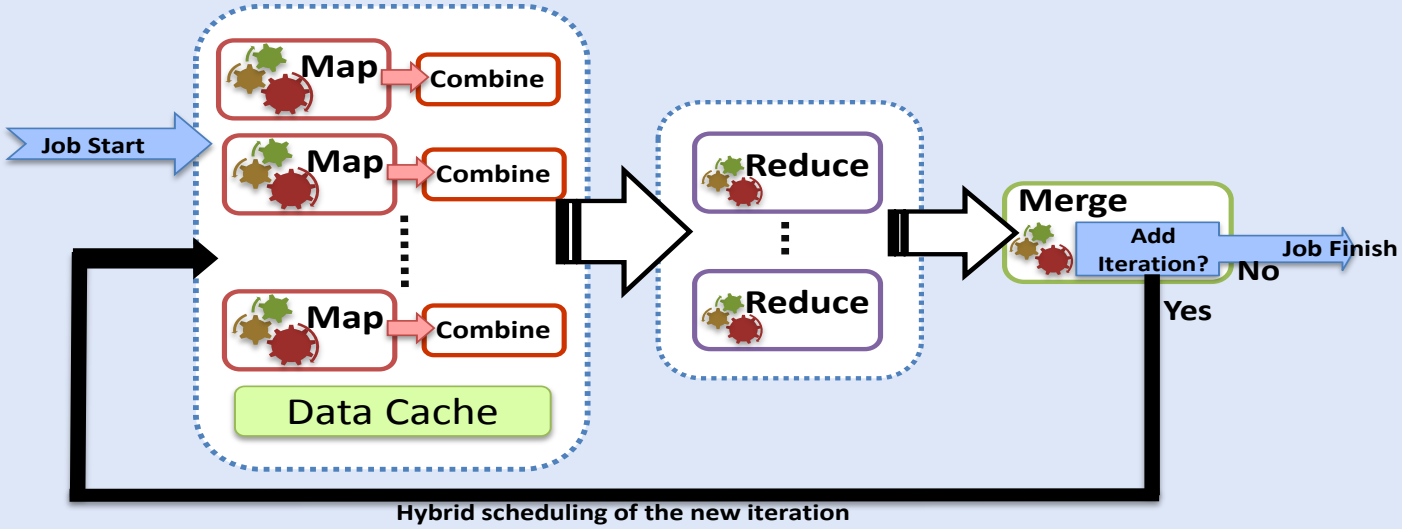

- Merge Step
- In-Memory Caching of static data
- Cache aware hybrid scheduling using Queues as well as using a bulletin board (special table)

Performance – Kmeans Clustering

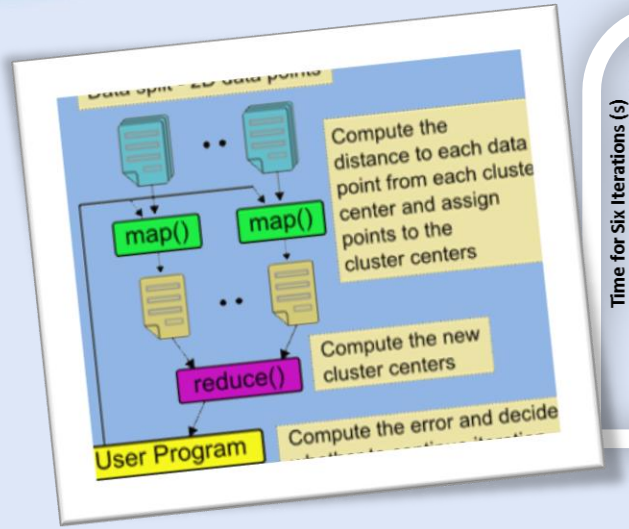

Performance with/without data caching

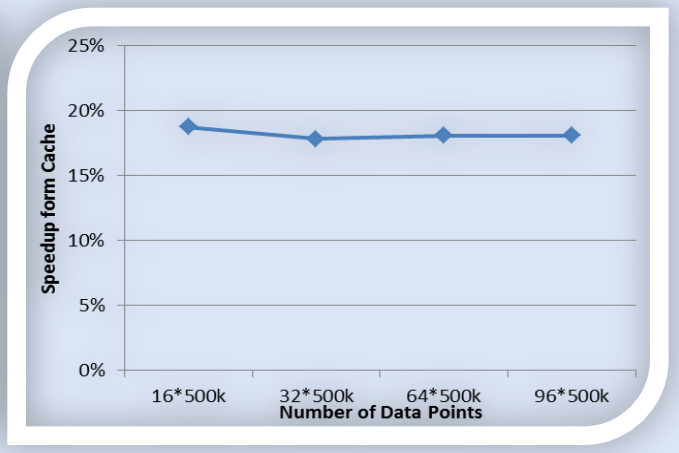

Speedup gained using data cache

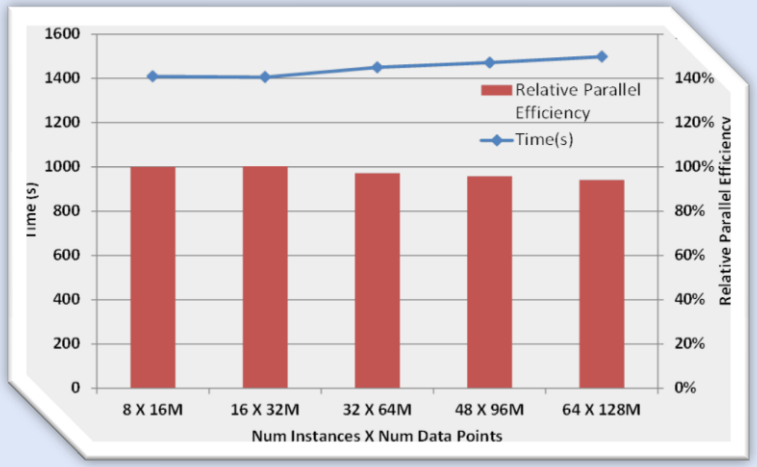

Scaling speedup

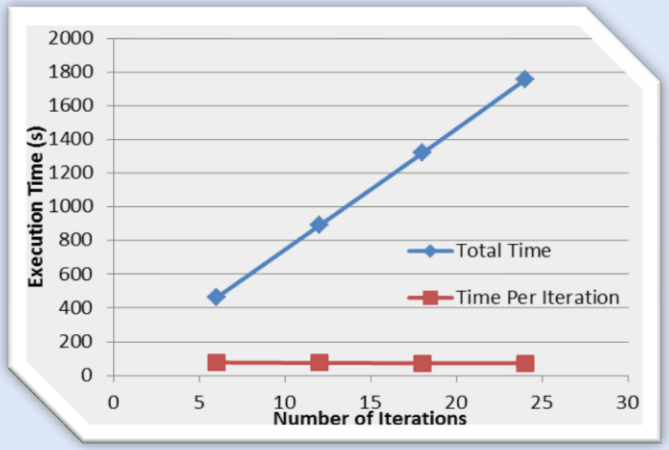

Increasing number of iterations

Performance Comparisons

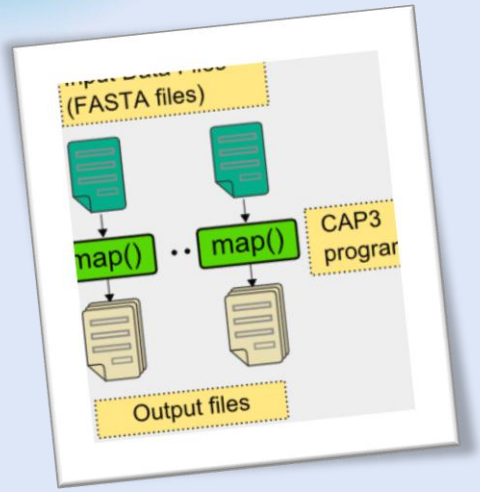

BLAST Sequence Search

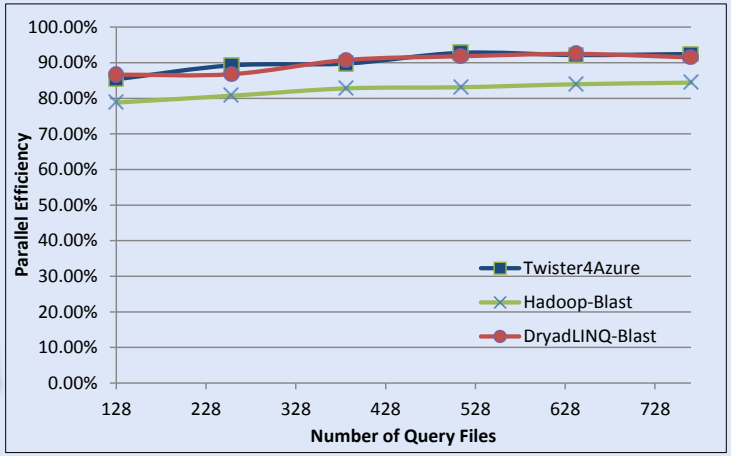

Smith Watermann Sequence Alignment

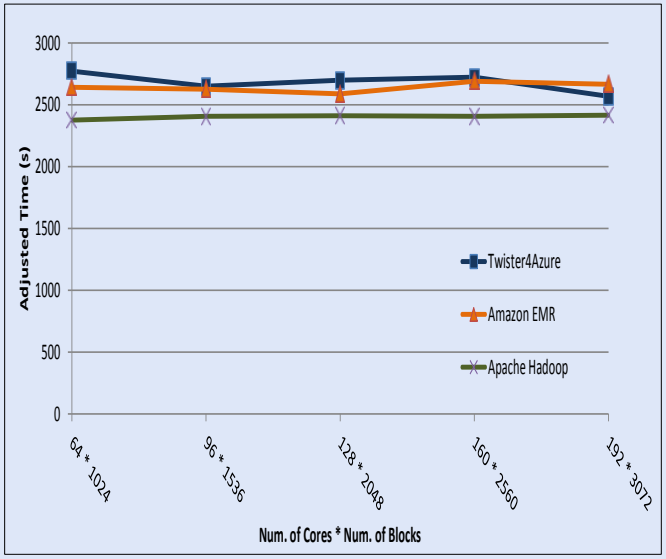

Cap3 Sequence Assembly

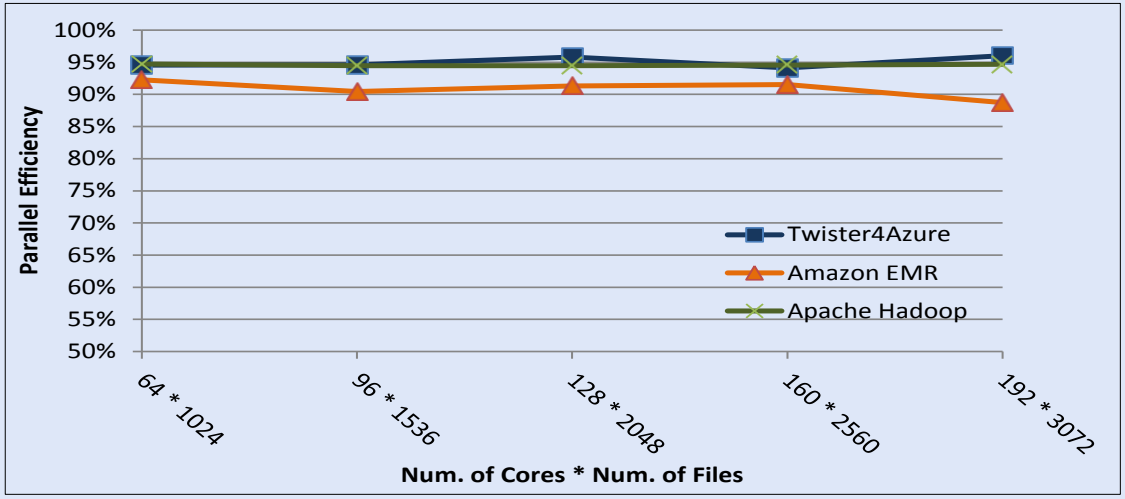

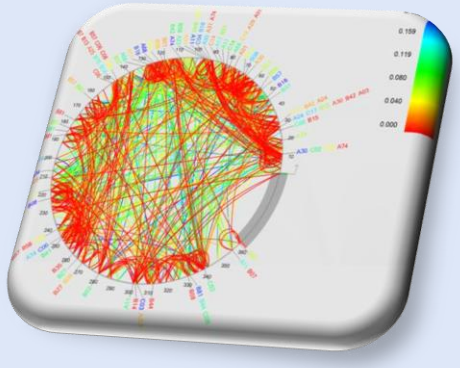

Twister as Map Collective

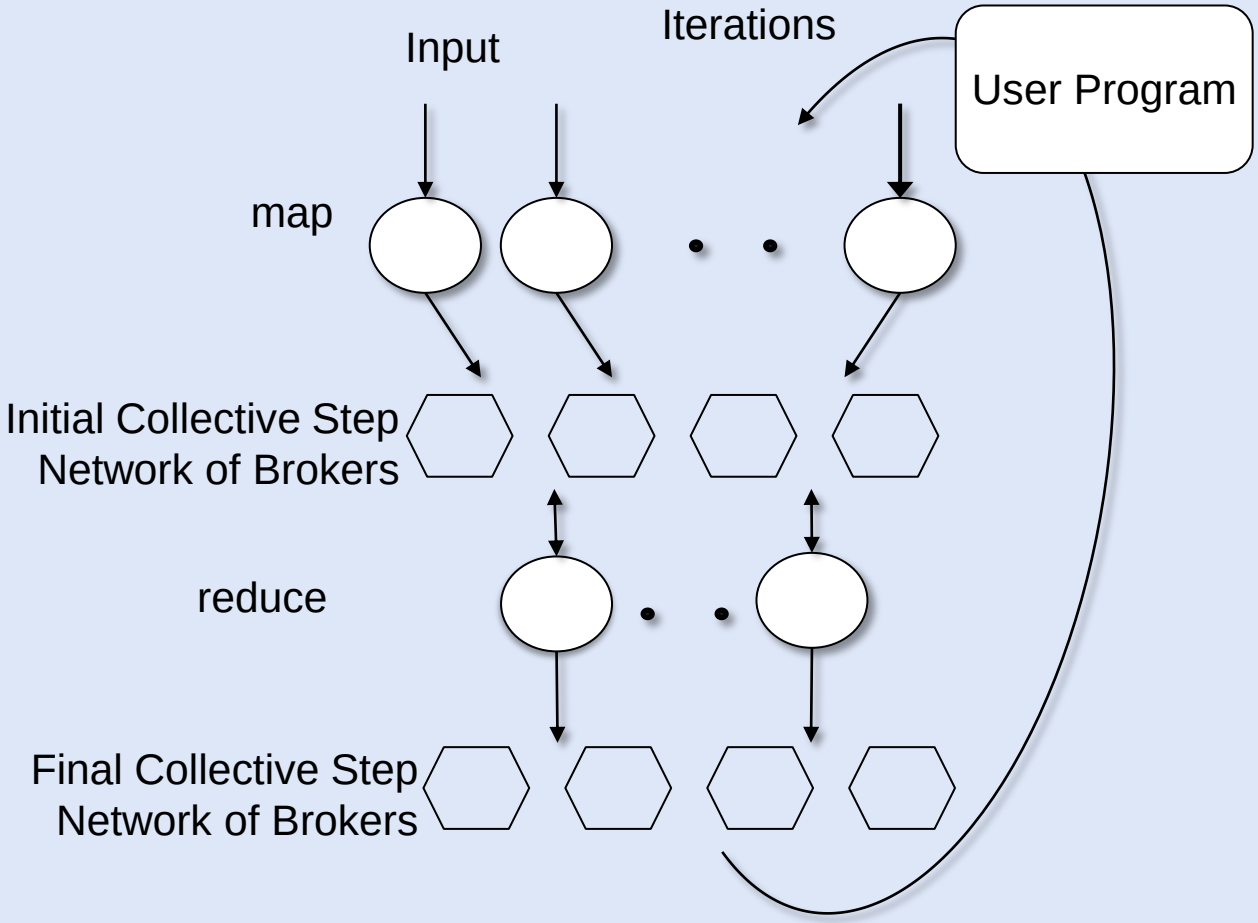

Usage Statistics from FutureGrid

- Based on User input we focused on
  - Nimbus (53%)
  - Eucalyptus (51%)
  - **Hadoop (37%)**
  - HPC (36%)

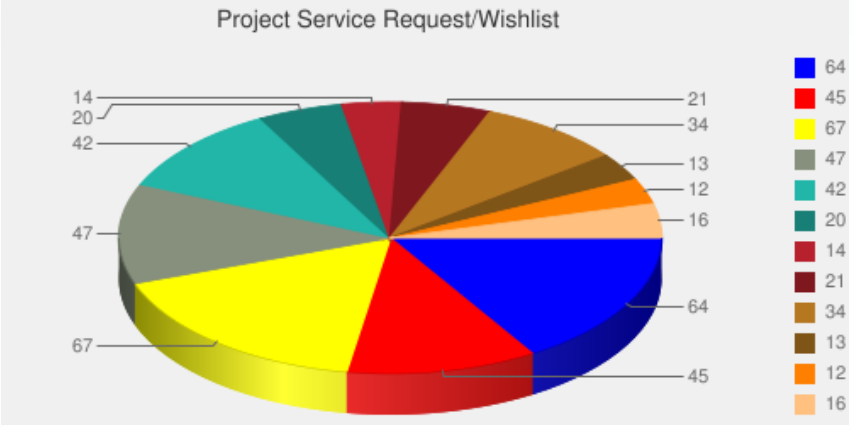

- Eucalyptus: 64(50.8%)
- High Performance Computing Environment: 45(35.7%)
- Nimbus: 67(53.2%)
- Hadoop: 47(37.3%)
- MapReduce: 42(33.3%)
- Twister: 20(15.9%)
- OpenNebula: 14(11.1%)
- Genesis II: 21(16.7%)
- Common TeraGrid Software Stack: 34(27%)
- Unicore 6: 13(10.3%)
- gLite: 12(9.5%)
- OpenStack: 16(12.7%)

Future Grid

Hadoop on FutureGrid

- **Goal**:
  - Simplify running Hadoop jobs thru FutureGrid batch queue systems
  - Allows user customized install of Hadoop

myHadoop

- **Status and Milestones**
  - **Today**
    - myHadoop 0.2a released early this year, deployed to Sierra and India, tutorial available
  - **In future**
    - deploy to Alamo, Hotel, Xray (end of year 2)

Future Grid

Many FutureGrid uses of MapReduce demonstrated in various Tutorials

- https://portal.futuregrid.org/tutorials
- Running Hadoop as a batch job using MyHadoop
  - Useful for coordinating many hadoop jobs through the HPC system and queues
- Running Hadoop on Eucalyptus
  - Running hadoop in a virtualized environment
- Running Hadoop on the Grid Appliance
  - Running haddop in a virtualized environment
  - Benefit from easy setup
- Eucalyptus and Twister on FG
  - Those wanting to use the Twister Iterative MapReduce
- Could organize workshops, seminars and/or increase online material.

Big Data for Science

300+ Students learning about Twister & Hadoop MapReduce technologies, supported by FutureGrid.

July 26-30, 2010 NCSA Summer School Workshop
http://salsahpc.indiana.edu/tutorial

ADMI Cloudy View on Computing Workshop June 2011

Concept and Delivery by Jerome Mitchell: Undergraduate ECSU, Masters Kansas, PhD (in progress)Indiana

- Jerome took two courses from IU in this area Fall 2010 and Spring 2011 on FutureGrid
- ADMI: Association of Computer and Information Science/Engineering Departments at Minority Institutions
- 10 Faculty and Graduate Students from ADMI Universities
- Included bootcamp on Mapreduce illustrated with case studies of scientific applications on FutureGrid.
- At the conclusion of the workshop, the participants indicated that they would incorporate cloud computing into their courses and/or research.

Future Grid

ADMI Cloudy View on Computing Workshop Participants

DeShea Simon
Hampton University

Timothy Holston
Mississippi Valley State University

Mohammad Hasan
Elizabeth City State University

Constance Bland
Mississippi Valley State University

Candace Adams
Auburn University

Felicia Doswell
Norfolk State University

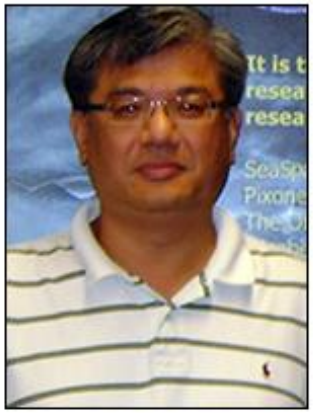

Yenhung Hu
Hampton University

Willie Fuller
Norfolk State University

Natarajan Meghanathan
Jackson State University

Darnell Johnson
Elizabeth City State University

Hadoop

- Hadoop provides an open source implementation of MapReduce and HDFS.
- myHadoop provides a set of scripts to configure and run Hadoop within an HPC environment
  - From San Diego Supercomputer Center
  - Available on India, Sierra, and Alamo systems within FutureGrid

myHadoop

- Log into to india & load mymadoop

```
user@host:$ ssh user@india.futuregrid.org

[user@i136 ~]$ module load myhadoop
myHadoop version 0.2a loaded
[user@i136 ~]$ echo $MY_HADOOP_HOME
/N/soft/myHadoop
```

myHadoop

- Create a PBS Job

```
#PBS -q batch
#PBS -N hadoop_job
#PBS -l nodes=4:ppn=1
#PBS -o hadoop_run.out
#PBS -e hadoop_run.err
#PBS -V

module load java
#### Set this to the directory where Hadoop configs should
   be generated
#...
export HADOOP_CONF_DIR="${HOME}/myHadoop-config"
```

myHadoop

```
#### Start the Hadoop cluster
echo "Start all Hadoop daemons"
$HADOOP_HOME/bin/start-all.sh
#$HADOOP_HOME/bin/hadoop dfsadmin -safemode leave

#### Run your jobs here
echo "Run some test Hadoop jobs"
$HADOOP_HOME/bin/hadoop --config $HADOOP_CONF_DIR dfs -mkdir Data
$HADOOP_HOME/bin/hadoop --config $HADOOP_CONF_DIR dfs -
copyFromLocal $MY_HADOOP_HOME/gutenberg Data
$HADOOP_HOME/bin/hadoop --config $HADOOP_CONF_DIR dfs -ls
Data/gutenberg
$HADOOP_HOME/bin/hadoop --config $HADOOP_CONF_DIR jar
$HADOOP_HOME/hadoop-0.20.2-examples.jar wordcount Data/gutenberg
Outputs
$HADOOP_HOME/bin/hadoop --config $HADOOP_CONF_DIR dfs -ls Outputs
$HADOOP_HOME/bin/hadoop --config $HADOOP_CONF_DIR dfs -
copyToLocal Outputs ${HOME}/Hadoop-Outputs
```

myHadoop

- Submit a job

```
[user@i136 ~]$ qsub pbs-example.sh
125525.i136
[user@i136 ~]$ qstat -u user

i136:

Req'd  Req'd   Elap
Job ID               Username Queue    Jobname          SessID NDS   TSK
Memory Time  S Time
-------------------- -------- -------- ---------------- ------ ----- ---
------ ----- - -----
125525.i136          user   batch    hadoop_job          --      4    4
--  04:00 Q   --
```

myHadoop

- Get results

```
[user@i136 ~]$ head Hadoop-Outputs/part-r-00000
"'After   1
"'My   1
"'Tis  2
"A  12
"About 2
"Ah!   2
"Ah!"  1
"Ah,   1
"All   2
"All!  1
```

Custom Hadoop

- Can use another configuration of Hadoop…

```
### Run the myHadoop environment script to set the
   appropriate variables
#
# Note: ensure that the variables are set correctly in
   bin/setenv.sh
. /N/soft/myHadoop/bin/setenv.sh
export HADOOP_HOME=${HOME}/my-custom-hadoop
```

Eucalyptus on FutureGrid

Slide authors:
Archit Kulshrestha, Gregor von Laszewski, Andrew Younge

http://www.youtube.com/watch?v=Cp-YzYlwPUg

Before you can use Eucalyptus

- Please make sure you have a portal account
  - https://portal.futuregrid.org
- Please make sure you are part of a valid FG project
  - You can either create a new one or
  - You can join an existing one with permission of the Lead
- Do not apply for an account before you have joined the project, your Eucalyptus account request will not be granted!

Eucalyptus

- Elastic Utility Computing Architecture Linking Your Programs To Useful Systems
  - Eucalyptus is an open-source software platform that implements IaaS-style cloud computing using the existing Linux-based infrastructure
  - IaaS Cloud Services providing atomic allocation for
    - Set of VMs
    - Set of Storage resources
    - Networking

Open Source Eucalyptus

- **Eucalyptus Features**
  - Amazon AWS Interface Compatibility
  - Web-based interface for cloud configuration and credential management.
  - Flexible Clustering and Availability Zones.
  - Network Management, Security Groups, Traffic Isolation
    - Elastic IPs, Group based firewalls etc.
  - Cloud Semantics and Self-Service Capability
    - Image registration and image attribute manipulation
  - Bucket-Based Storage Abstraction (S3-Compatible)
  - Block-Based Storage Abstraction (EBS-Compatible)
  - Xen and KVM Hypervisor Support

Source: http://www.eucalyptus.com

Eucalyptus Testbed

- Eucalyptus is available to FutureGrid Users on the India and Sierra clusters.
- Users can make use of a maximum of 50 nodes on India. Each node supports up to 8 small VMs. Different Availability zones provide VMs with different compute and memory capacities.

```
AVAILABILITYZONE        india   149.165.146.135
AVAILABILITYZONE        |- vm types     free / max   cpu   ram  disk
AVAILABILITYZONE        |- m1.small     0400 / 0400   1    512     5
AVAILABILITYZONE        |- c1.medium    0400 / 0400   1   1024     7
AVAILABILITYZONE        |- m1.large     0200 / 0200   2   6000    10
AVAILABILITYZONE        |- m1.xlarge    0100 / 0100   2   12000    10
AVAILABILITYZONE        |- c1.xlarge    0050 / 0050   8   20000    10
```

Eucalyptus Account Creation

- Use the Eucalyptus Web Interfaces at

  https://eucalyptus.india.futuregrid.org:8443/

- On the Login page click on Apply for account.
- On the next page that pops up fill out ALL the Mandatory AND optional fields of the form.
- Once complete click on signup and the Eucalyptus administrator will be notified of the account request.
- You will get an email once the account has been approved.
- Click on the link provided in the email to confirm and complete the account creation process.

Obtaining Credentials

- Download your credentials as a zip file from the web interface for use with euca2ools.
- Save this file and extract it for local use or copy it to India/Sierra.
- On the command prompt change to the euca2-{username}-x509 folder which was just created.
  - cd euca2-username-x509
- Source the eucarc file using the command source eucarc.
  - source ./eucarc

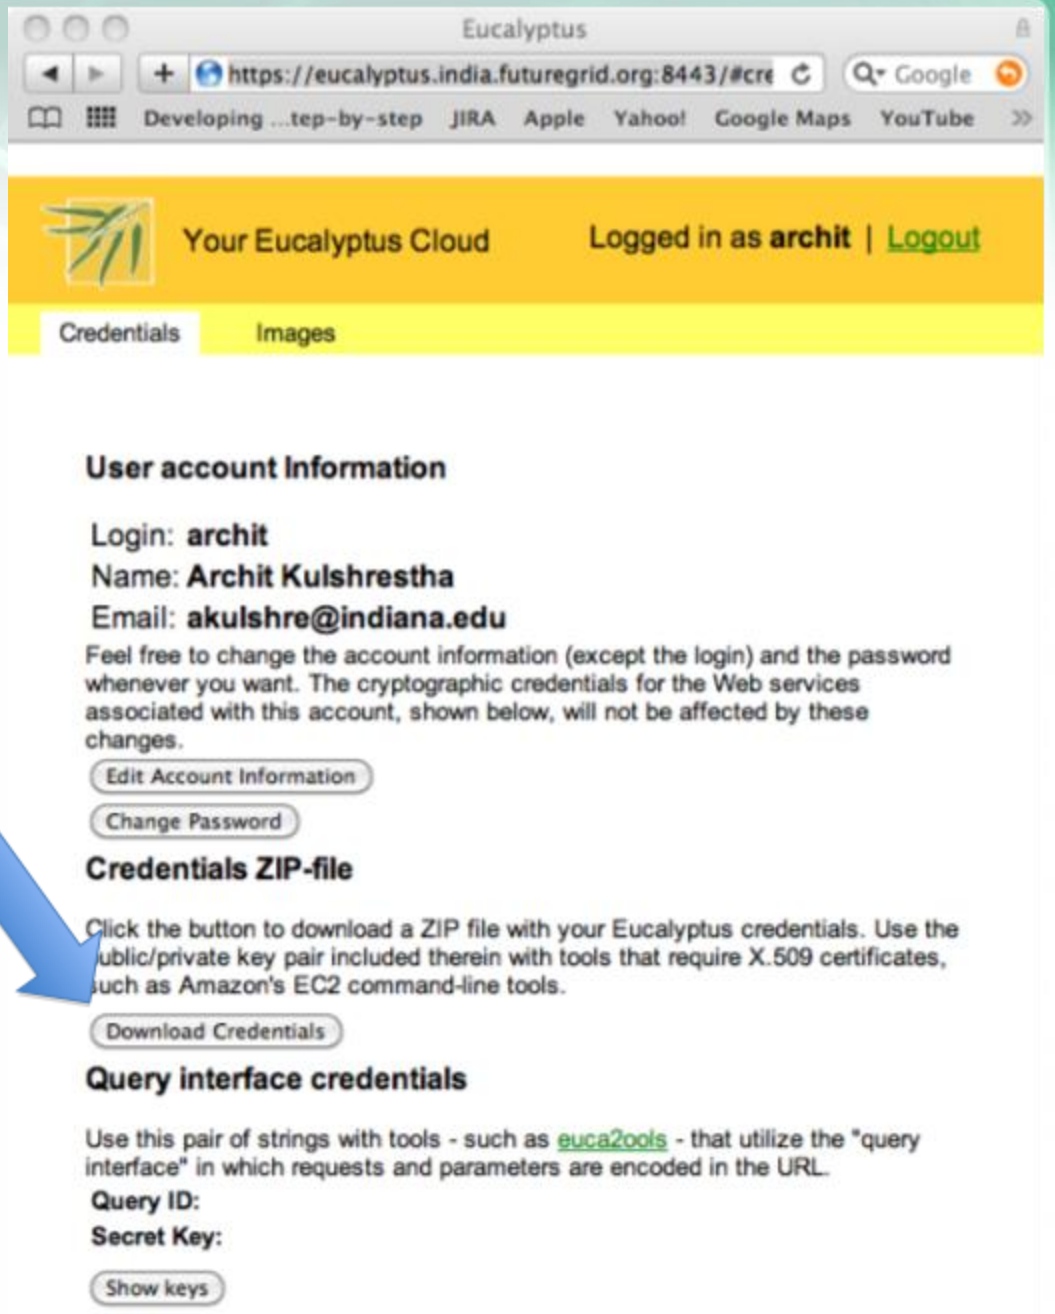

Install/Load Euca2ools

- Euca2ools are the command line clients used to interact with Eucalyptus.
- If using your own platform Install euca2ools bundle from http://open.eucalyptus.com/downloads
  - Instructions for various Linux platforms are available on the download page.
- On FutureGrid log on to India/Sierra and load the Euca2ools module.

```
$ module load euca2ools
euca2ools version 1.2 loaded
```

Euca2ools

- Testing your setup
  - Use euca-describe-availability-zones to test the setup.
- List the existing images using euca-describe-images

```
euca-describe-availability-zones
AVAILABILITYZONE india 149.165.146.135
```

```
$ euca-describe-images
IMAGE emi-0B951139 centos53/centos.5-3.x86-64.img.manifest.xml admin
available public x86_64 machine
IMAGE emi-409D0D73 rhel55/rhel55.img.manifest.xml admin available public
x86_64 machine
…
```

Key management

- Create a keypair and add the public key to eucalyptus.

```
$ euca-add-keypair userkey > userkey.pem
```

- Fix the permissions on the generated private key.

```
$ chmod 0600 userkey.pem
```

```
$ euca-describe-keypairs
KEYPAIR userkey
    0d:d8:7c:2c:bd:85:af:7e:ad:8d:09:b8:ff:b0:54:d5:8c:66:86:5d
```

Image Deployment

- Now we are ready to start a VM using one of the pre-existing images.
- We need the emi-id of the image that we wish to start. This was listed in the output of euca-describe-images command that we saw earlier.
  - We use the euca-run-instances command to start the VM.

```
$ euca-run-instances -k userkey -n 1 emi-0B951139 -t c1.medium
RESERVATION r-4E730969 archit archit-default
INSTANCE i-4FC40839 emi-0B951139 0.0.0.0 0.0.0.0 pending userkey 2010-07-20T20:35:47.015Z eki-78EF12D2 eri-5BB61255
```

Monitoring

- euca-describe-instances shows the status of the VMs.

```
$ euca-describe-instances
RESERVATION r-4E730969 archit default
INSTANCE i-4FC40839 emi-0B951139 149.165.146.153 10.0.2.194 pending
userkey 0 m1.small 2010-07-20T20:35:47.015Z india eki-78EF12D2 eri-
5BB61255
```

- Shortly after…

```
$ euca-describe-instances
RESERVATION r-4E730969 archit default
INSTANCE i-4FC40839 emi-0B951139 149.165.146.153 10.0.2.194 running
userkey 0 m1.small 2010-07-20T20:35:47.015Z india eki-78EF12D2 eri-
5BB61255
```

VM Access

- First we must create rules to allow access to the VM over ssh.

```
euca-authorize -P tcp -p 22 -s 0.0.0.0/0 default
```

- The ssh private key that was generated earlier can now be used to login to the VM.

```
ssh -i userkey.pem root@149.165.146.153
```

Image Deployment (1/3)

- We will use the example Fedora 10 image to test uploading images.
  - Download the gzipped tar ball

```
wget http://open.eucalyptus.com/sites/all/modules/pubdlcnt/pubdlcnt.php?file=http://www.eucalyptussoftware.com/downloads/eucalyptus-images/euca-fedora-10-x86_64.tar.gz&amp;amp;nid=1210
```

- Uncompress and Untar the archive

```
tar zxf euca-fedora-10-x86_64.tar.gz
```

Image Deployment (2/3)

- Next we bundle the image with a kernel and a ramdisk using the euca-bundle-image command.
  - We will use the xen kernel already registered.
    - euca-describe-images returns the kernel and ramdisk IDs that we need.

```
$ euca-bundle-image -i euca-fedora-10-x86_64/fedora.10.x86-64.img --
    kernel eki-78EF12D2 --ramdisk eri-5BB61255
```

- Use the generated manifest file to upload the image to Walrus

```
$ euca-upload-bundle -b fedora-image-bucket -m /tmp/fedora.10.x86-
    64.img.manifest.xml
```

Image Deployment (3/3)

- Register the image with Eucalyptus

```
euca-register fedora-image-bucket/fedora.10.x86-64.img.manifest.xml
```

- This returns the image ID which can also be seen using euca-describe-images

```
$ euca-describe-images
IMAGE emi-FFC3154F fedora-image-bucket/fedora.10.x86-
    64.img.manifest.xml archit available public x86_64 machine eri-
    5BB61255 eki-78EF12D2
IMAGE emi-0B951139 centos53/centos.5-3.x86-64.img.manifest.xml
    admin available public x86_64 machine ...
```