Empowering Data Management, Diagnosis, and Visualization of Cloud-Resolving Models (CRM) by Cloud Library upon Spark and Hadoop

Wei-Kuo TAO (GSFC)
Xian-He Sun (IIT)
Shujia Zhou (Northrop Grumman Information Technology)
Toshihisa Matsui (ESSIC UMD)
Xiaowen Li (GESTAR MSU)
Dan Duffy (Collaborator, GSFC)

ESTO 2016 Meeting
June 15, 2016

Earth Science Technology Office

Goals

- Make Cloud Resolving Model output more usable by science community
  - Accelerate visualization of output.
  - Inter-compare large volumes of output from high-resolution simulations.
  - Diagnose key processes for cloud-precipitation.

- Demonstrate the value to distribute, visualize, analyze and inter-compare Cloud Resolving Model output and data with GCE and NU-WRF

GCE: Goddard Cumulus Ensemble model (1982 -)
NU-WRF: NASA Unified Weather Research Forecast (2010 -)

Approach

- Develop Super Cloud Library (SCL) supporting Cloud Resolving Model Data Analyses using Spark on Hadoop.
  - *Create cloud data files;*
  - *Develop data model and Hadoop format transformer;*
  - *Develop a dynamic Hadoop reader tool;*
  - *Develop subsetting and visualization APIs (Application Programming Interfaces);*
  - *Develop a Web User Interface.*
- Conduct Demo of GCE and NU-WRF diagnoses on NCCS.

Super Cloud Library (SCL)

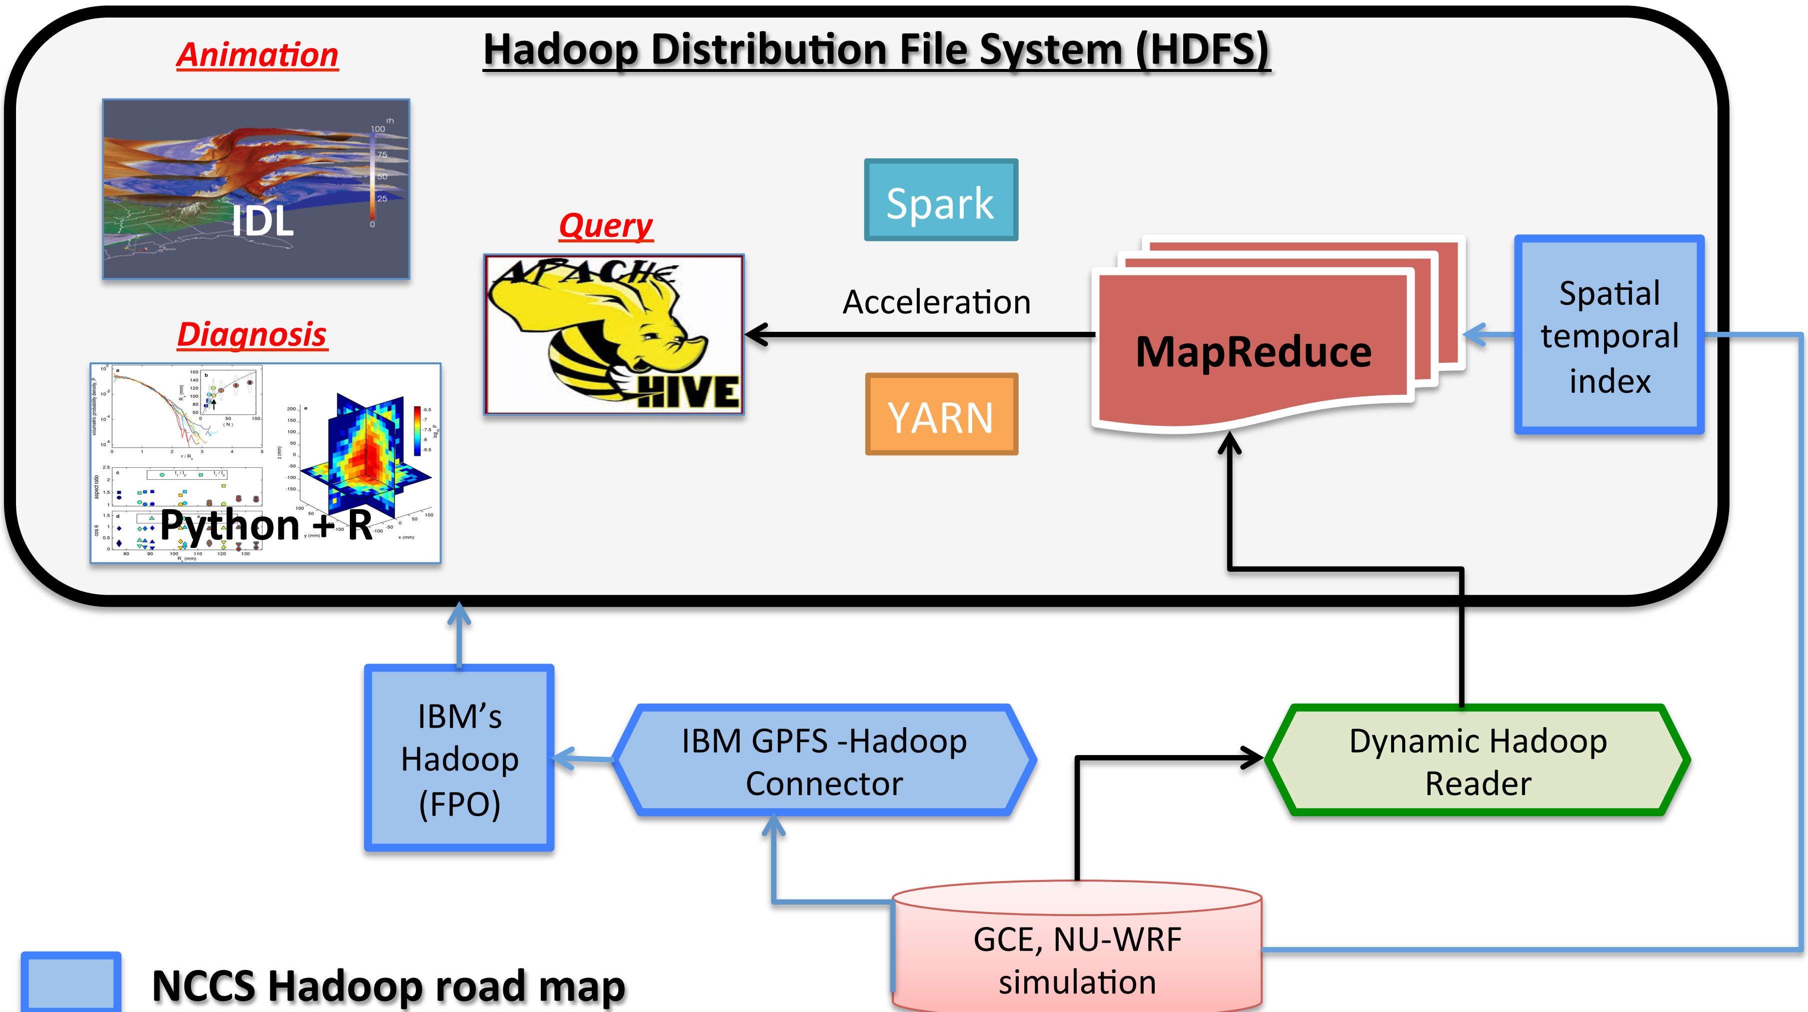

Outline

- Cloud-Model Data (Real Cases);
- Hadoop Format Translator;
- Subsetting/Visualization via Hadoop-IDL, Hardoop-R and Model Inter-Comparison/Diagnoses;
- Hadoop Dynamic Reader: GPFS and HIVE Interface;
- SCL Web Design;
- Future Work.

NU-WRF Real Cases

Long-Term Case

- **Grid (9km)**: 600x400x50
- **Date**: West African Monsoon (June-July-August in 2006)
- **Output frequency**: 3hr
- **Data Sizes**: 0.34TB

Semi-Giga Cases

- **Grid (2km)**: 2500x2500x50
- **Date**: Tornado Outbreak (6days), Tropical Storm Bill (6days)
- **Output frequency**: 1hr
- **Data Sizes**: 145 files x 2, 1.73TB x 2

NU-WRF Semi-Giga Cases (Δ=2km)

-2014 Tornado Outbreak vs Hurrinca Bill-

Total grid points (3D) is 0.312 Giga points

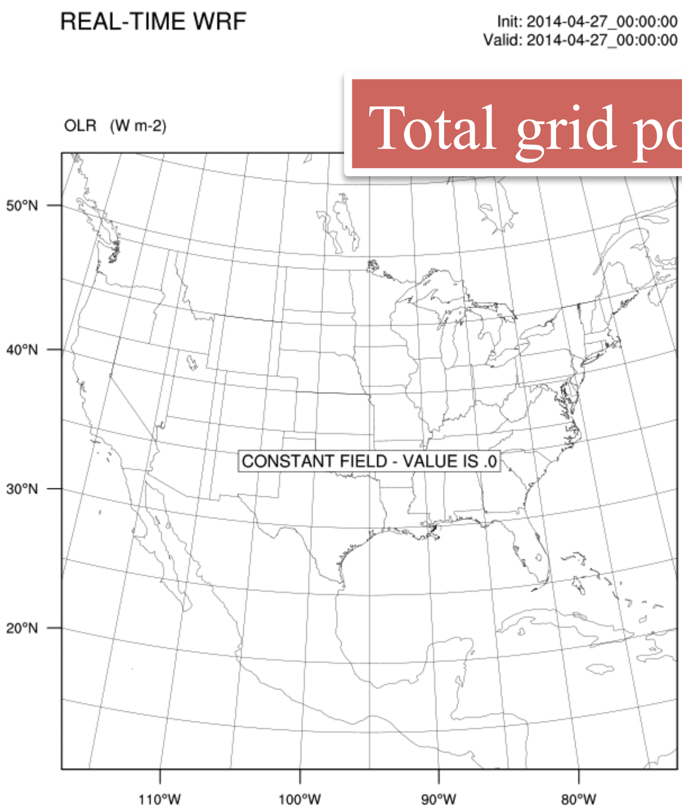

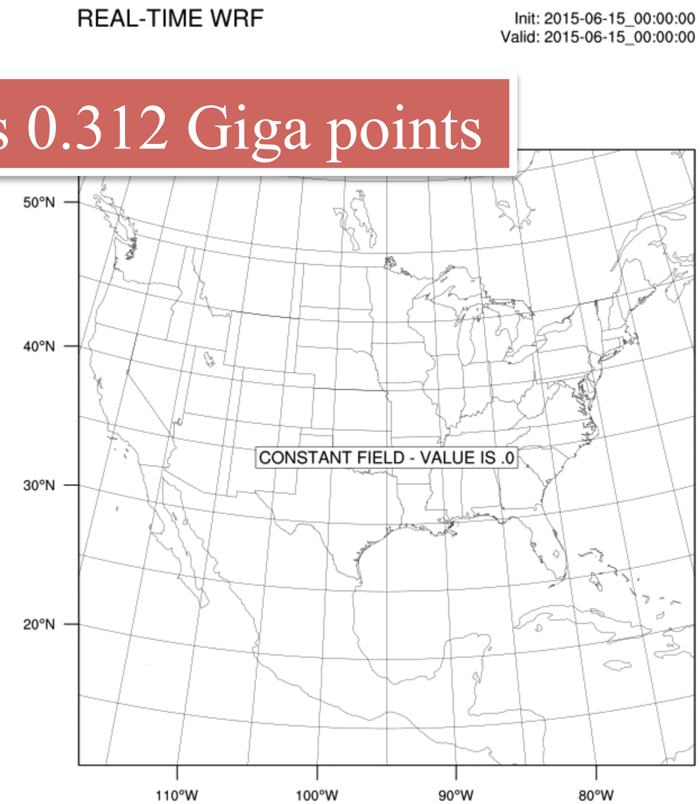

2500x2500x50

2014 Tornado Outbreak

2015 Tropical Storm Bill

OUTPUT FROM WRF V3.5.1 MODEL
WE = 2501 ; SN = 2501 ; Levels = 51 ; Dis = 2km ; Phys Opt = 56 ; PBL Opt = 5 ; Cu Opt = 0

OUTPUT FROM WRF V3.5.1 MODEL
WE = 2501 ; SN = 2501 ; Levels = 51 ; Dis = 2km ; Phys Opt = 56 ; PBL Opt = 5 ; Cu Opt = 0

Earth Science Technology Office

GCE Real Cases

1.8 Billion
3D grids

Long-Term MJO Case	**Giga Case**
• **Grid (1km):** 1024x1024x45	• **Grid (250m):** **4096x4096x106**
• **Case & Date:** DYNAMO (Nov. 1~Dec. 10, 2011)	• **Case & Date:** DYNAMO (Nov. 23~Nov. 29, 2011)
• **Output frequency:** 3hr	• **Output frequency:** 1hr
• **Data Sizes:** 0.832TB	• **Data Sizes:** 15 TB

GCE Giga Case (Δ=0.25km)
DYNAMO simulated vertical velocity details

4096x4096x106 grid points

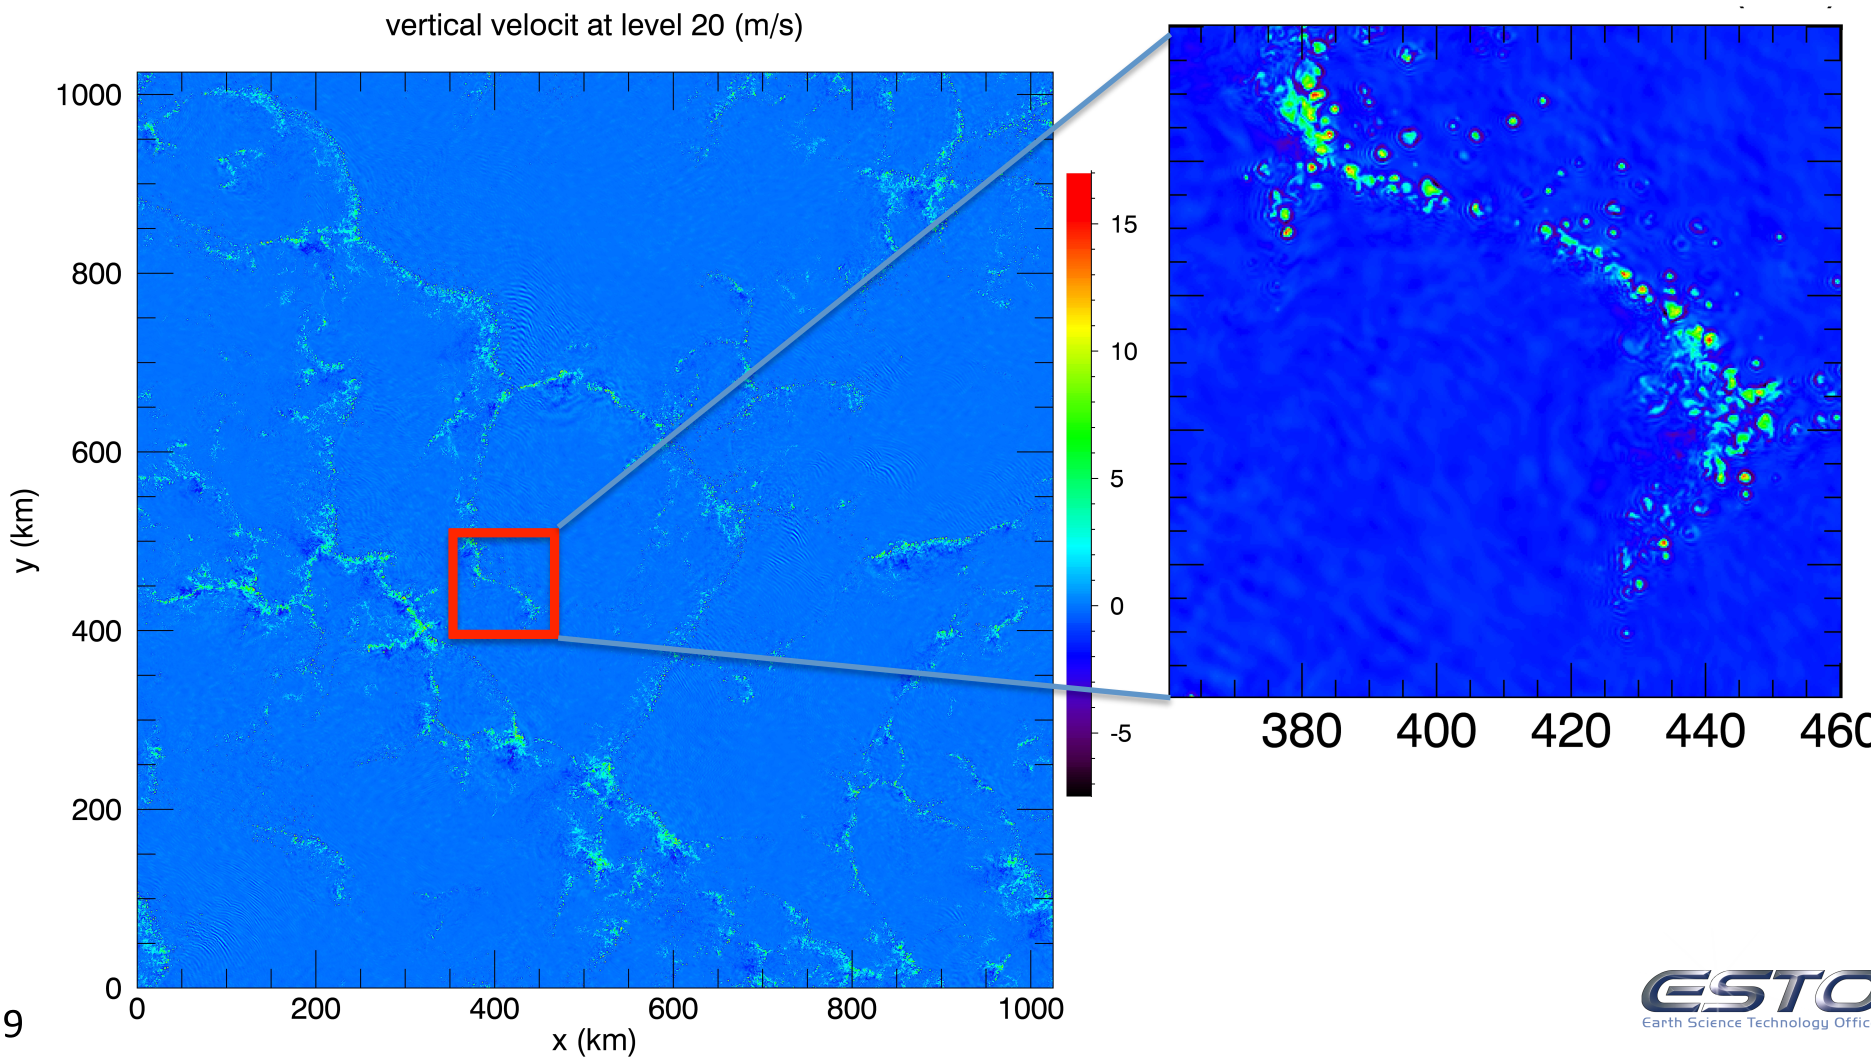

Summary: Data Model

	GCE		NU-WRF		
Case Tag	Long-Term (DYNAMO)	Giga Scale (DYNAMO)	Long-Term (AMMA)	Semi-Giga scale (Tornado 2014)	Semi-Giga Scale (Tropical Storm Bill 2015)
Grid points (i-j-k)	1024x1024x45	4096x4096x106	600x400x50	2500x2500x50	
Horizontal Grid	1km	0.25km	9km	2km	
Date	11/01/2011 - 12/10/2011	11/23/2011 - 11/29/2011	06/01/2006 - 09/01/2006	04/27/2014 - 05/03/2014	06/15/2015 – 06/21/2015
Model Integration (Output freq.)	5 weeks (3hr)	6 days (1hr)	3 months (3hr)	6 days (1hr)	
Total native output size	0.832TB	15TB	0.74TB	1.73TB	
Status of Porting HDFS	Yes	No[1]	Yes	Yes	Yes

Summary of SCL data model, including simulations of Goddard Cumulus Ensemble (GCE) model and NASA-Unified Weather Research and Forecasting (NU-WRF) model. 1. Our current quota of HDFS is not enough of porting Giga-scale GCE output into NCCS's HDFS yet, but in progress.

Hadoop Format Translator

Earth Science Technology Office

CSV Data-Size Reduction

CSV format (V1.0)

- Original CSV: each files has time-geolocation information and single geophysical parameter..

Geolocation+Parameter File x 27

ID,time,lev,lat,lon,param

```
1,2820,1,1.45085,75.82308,0.10038067E+04
2,2820,1,1.45090,75.83212,0.10038067E+04
3,2820,1,1.45095,75.84116,0.10038067E+04
4,2820,1,1.45100,75.85021,0.10038067E+04
……
```

HDFS

CSV format (V2.1)

- Geolocation information is stored in **single** separate file per case.
- Parameter files contains several parameters.
- **Parallel gzip program** compress CSV files quickly (4days → 7min).
- Hadoop reads the geolocation file and parameter files simultaneously.

Geolocation files

ID,i,j,lev,lat,lon

```
1,2820,1,1.45085,75.82308
2,2820,1,1.45090,75.83212
3,2820,1,1.45095,75.84116
4,2820,1,1.45100,75.85021
……
```

Parameter files x 3

ID, time, param1, param2…

```
1,0.10038067E+04, …
2,0.10038067E+04, …
3,0.10038067E+04, …
4,0.10038067E+04, …
……
```

PGZIP.sh

HDFS

This method is close to what NetCDF reader does.

CFMC CSV Data Size

Single Time Frame of NU-WRF Semi-Giga Case

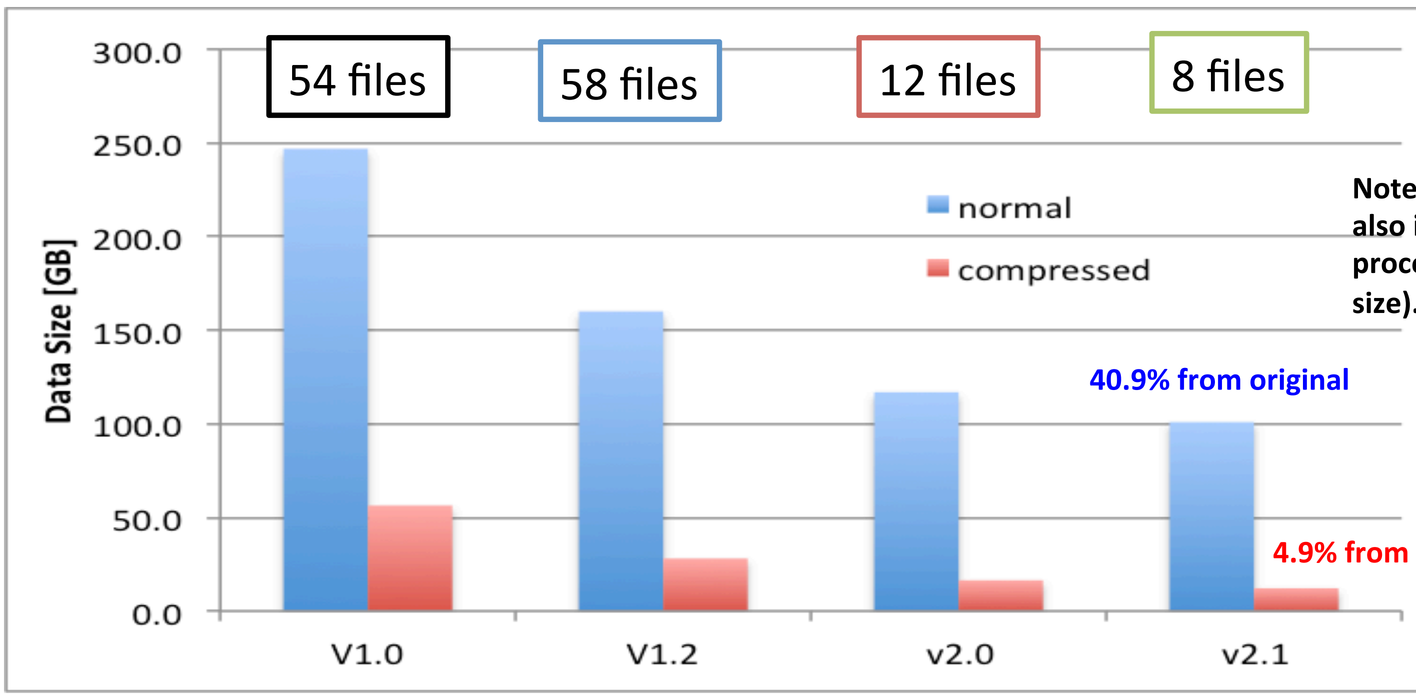

Note. Data size reduction also improved conversion process time (linearly in size).

40.9% from original

4.9% from original

V1.0: Original format

V1.2: Separate Geolocation from parameter files

V2.0: Bundle multiple parameters

V2.1: Only single Geolocation file for each case.

Subset/Visualization via Hadoop-IDL, Hadoop-R

Model Inter-comparison (Diagnosis)

Visualization with **IDL+Hadoop** and Subsetting with Impala: GCE 1024x1024x45 Updraft Simulation

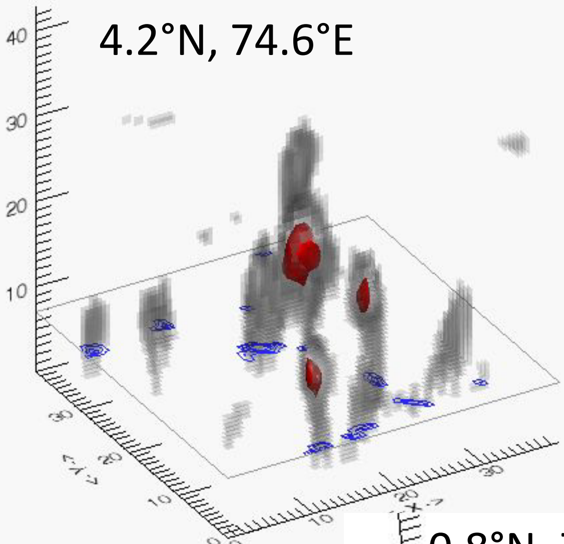

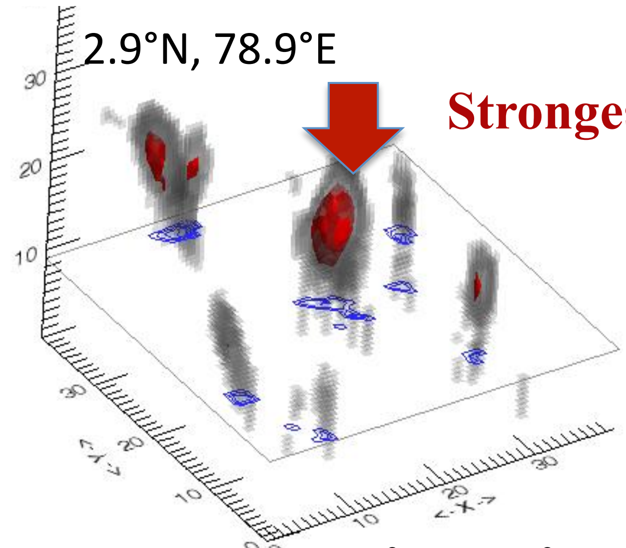

Strongest updraft

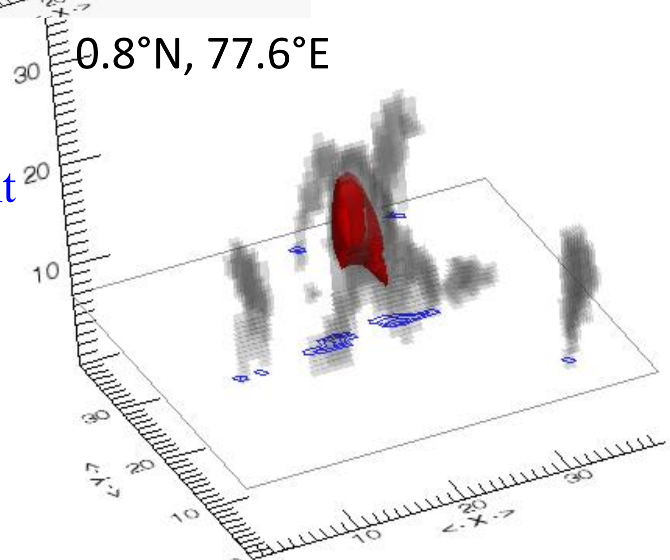

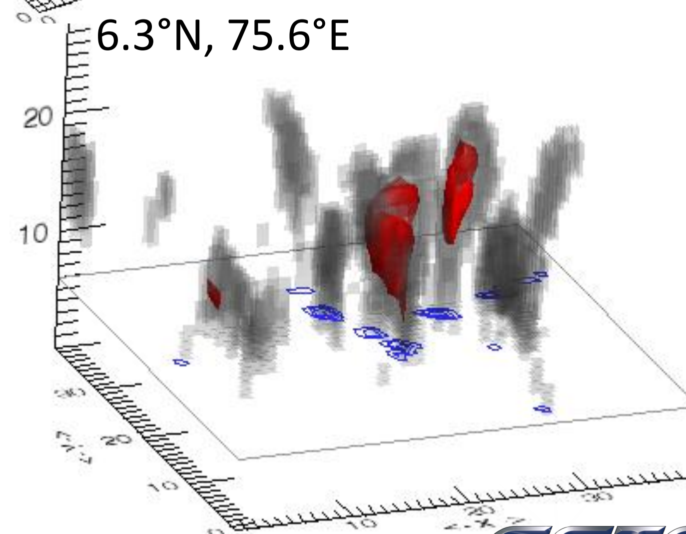

Four examples for subsetting at different times

Red color indicating the strongest updraft cores

Impala Subset Can Overlay Multiple Variables and Study Their Relations

GCE Simulation
1024 x 1024 x 45 Grid

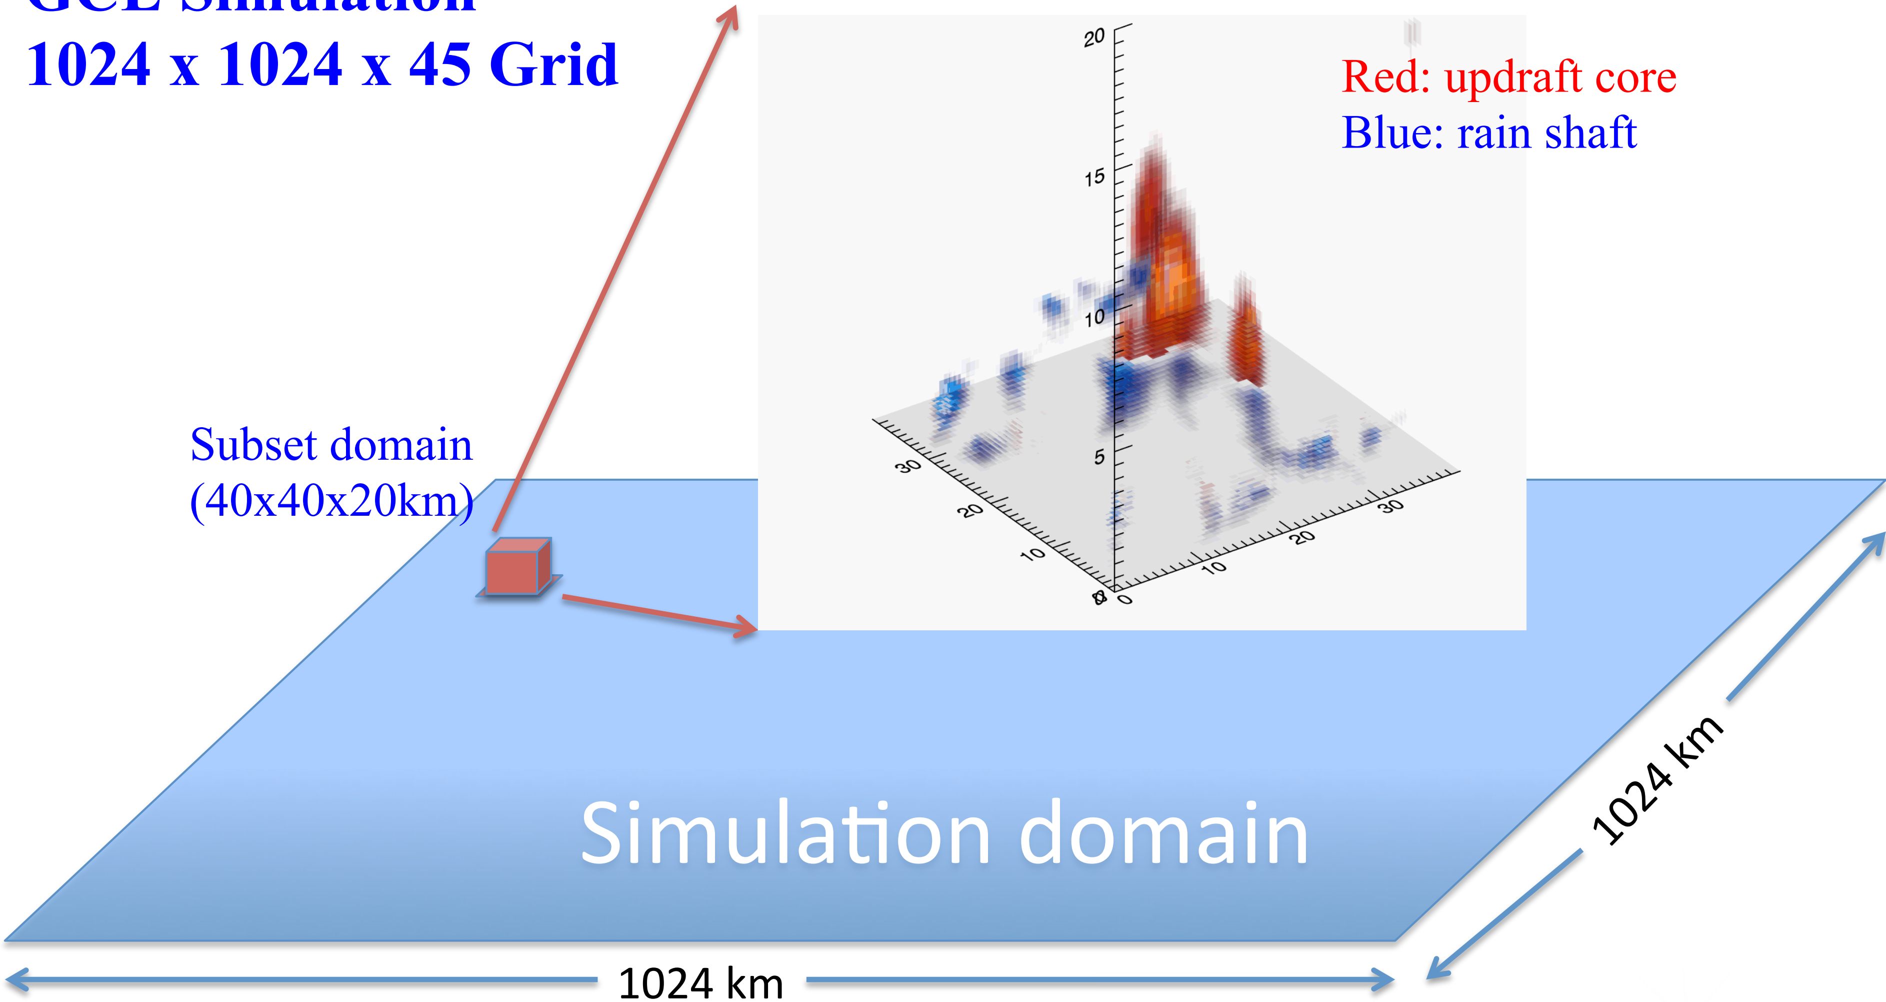

WRF Semi-giga Scale (2500x2500) Simulation

Surface Rainfall Rate with Impala Subsetting of Maximum w in Red Box

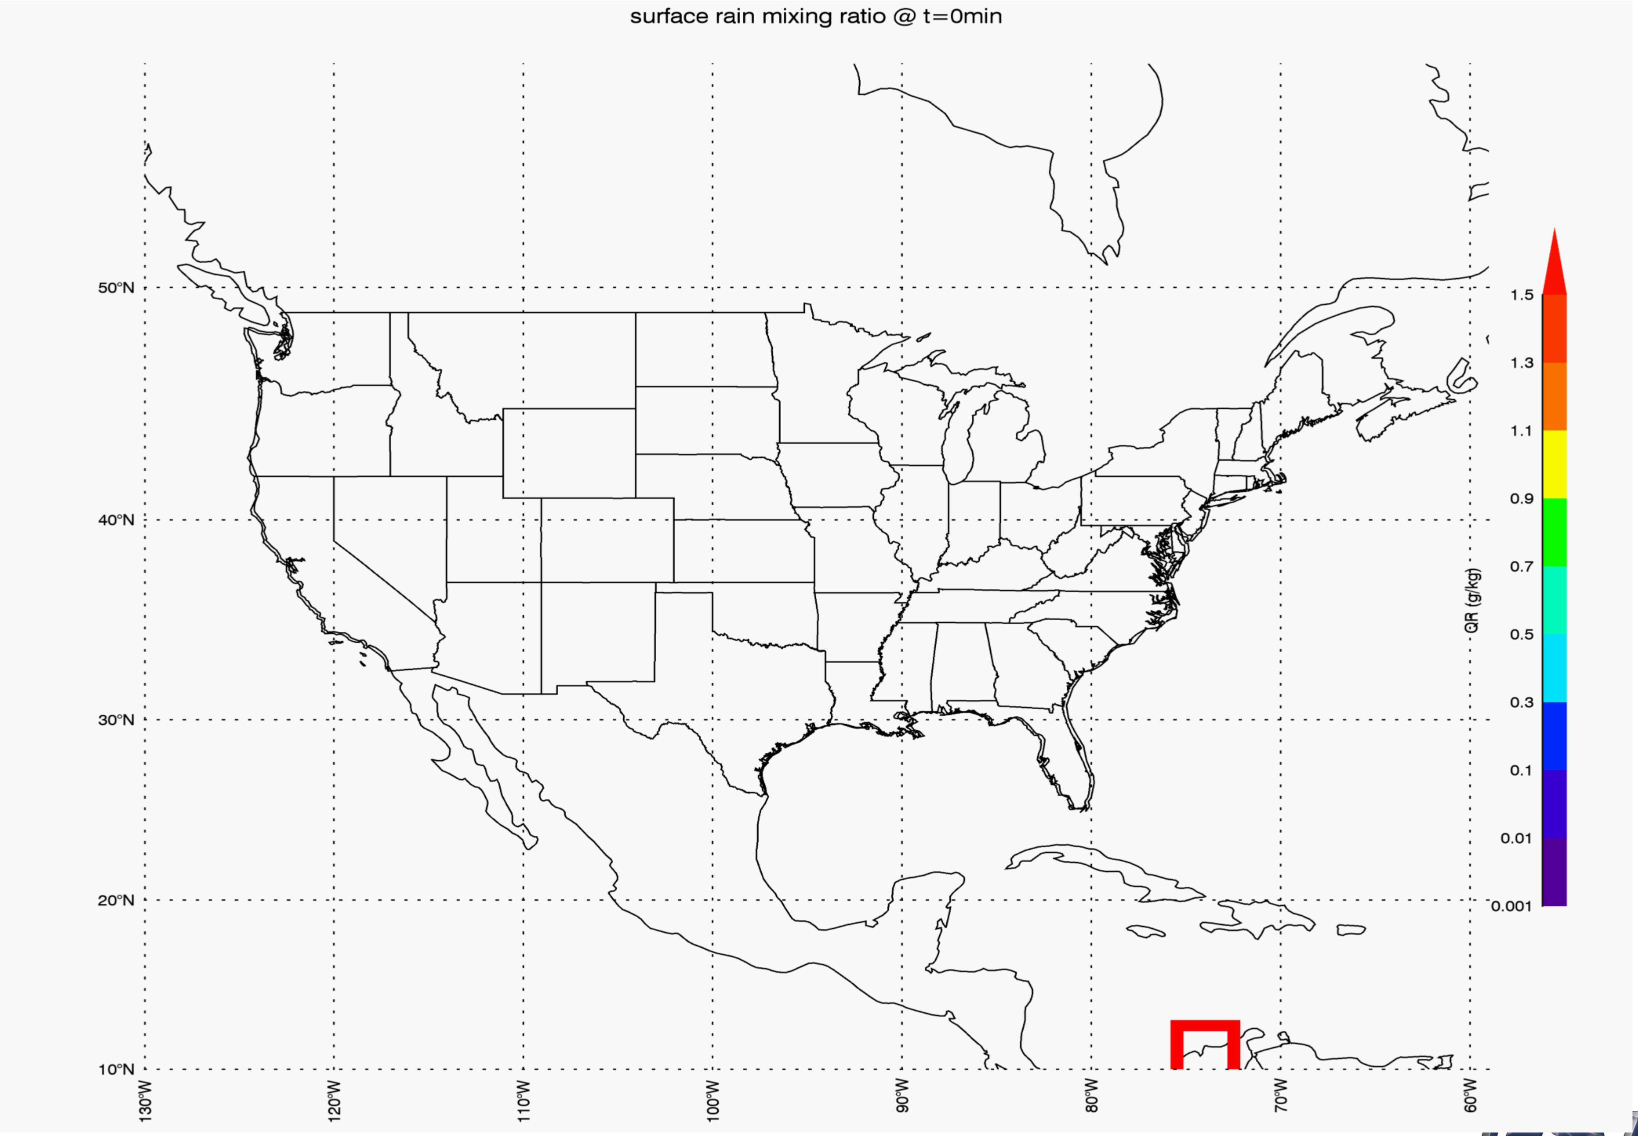

Earth Science Technology Office

3-D animation of subset data

Cloud ice + Rain

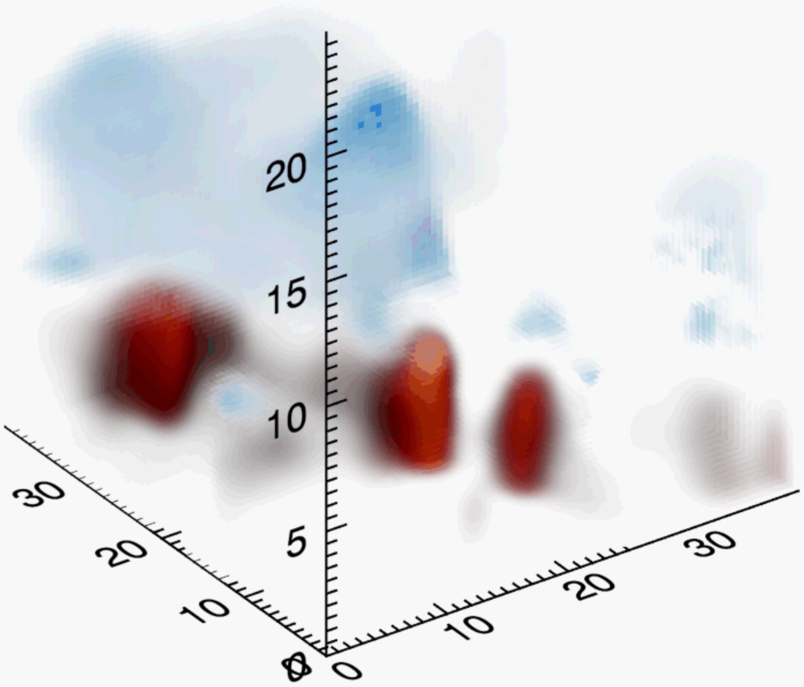

Downdraft + Updraft

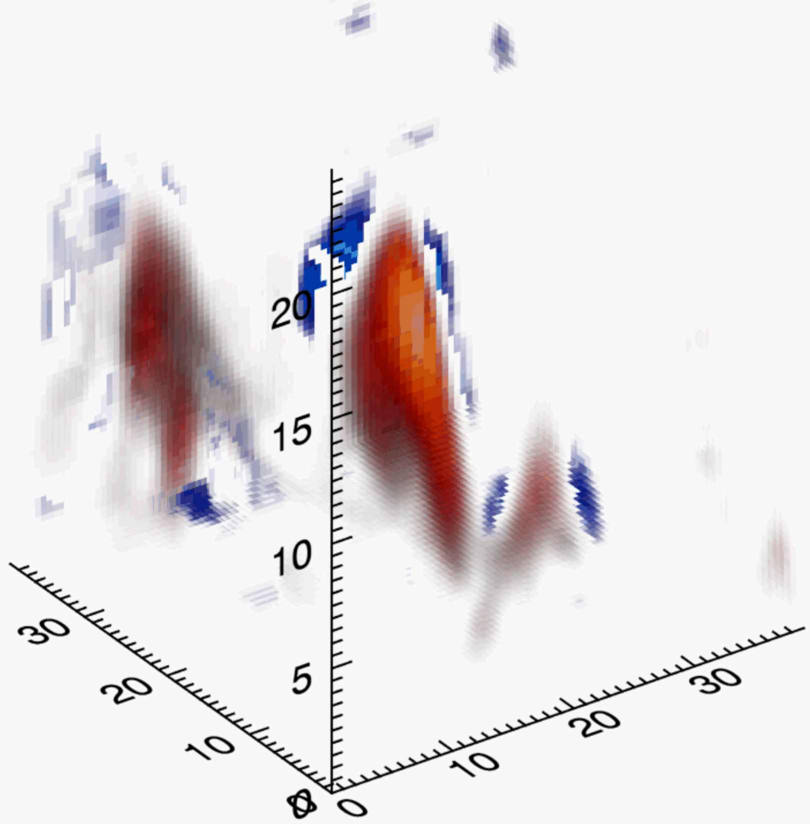

Spark-Python CRM Diagnostic Module

Model Inter-comparisons

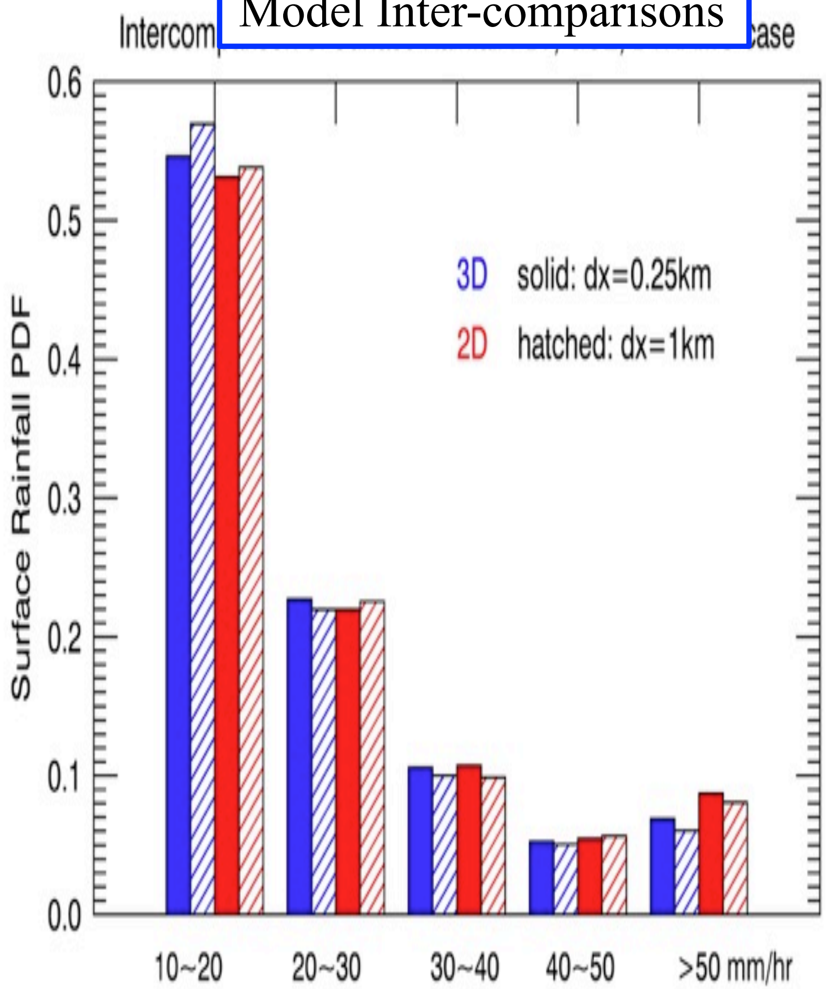

Scientific Diagnoses

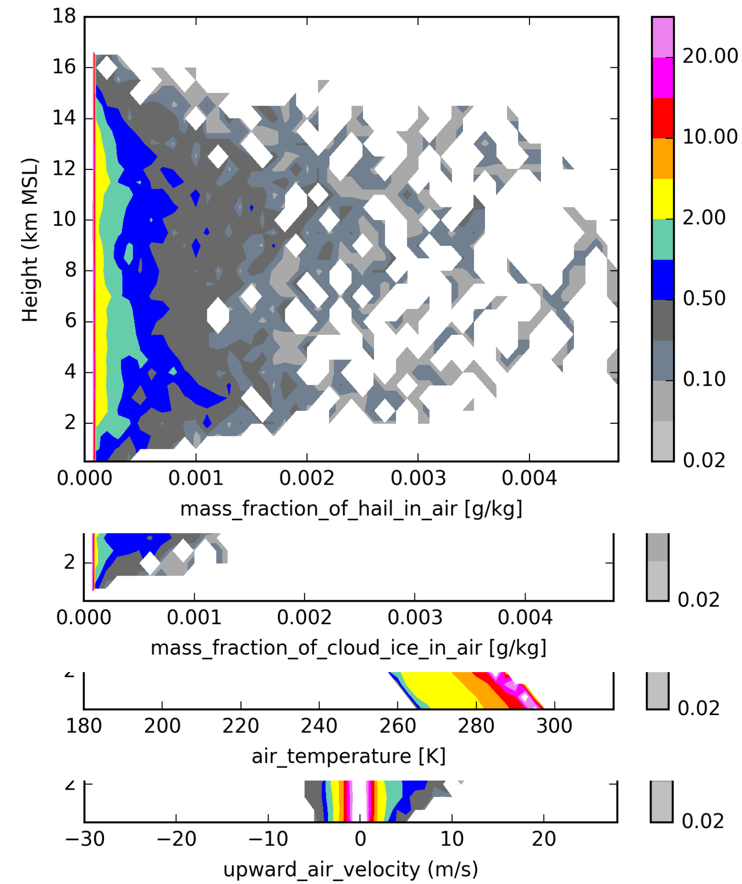

Earth Science Technology Office

Summary

- Different approaches for data subsetting and visualization have been experimented: e.g., Hadoop/Hive + IDL, Hadoop/Impala + IDL, Hadoop + R;
- Lesson learned:
  - Subsetting with Impala is much faster than Hive;
- Spark – Python is the current choice of model diagnoses and inter-comparisons.

ESTO Earth Science Technology Office

Dynamic Hadoop Reader and Visualization with R, Hadoop, Spark and Adaptive Subsetting of Earth Science Data in HDFS

Super Cloud Library (SCL)

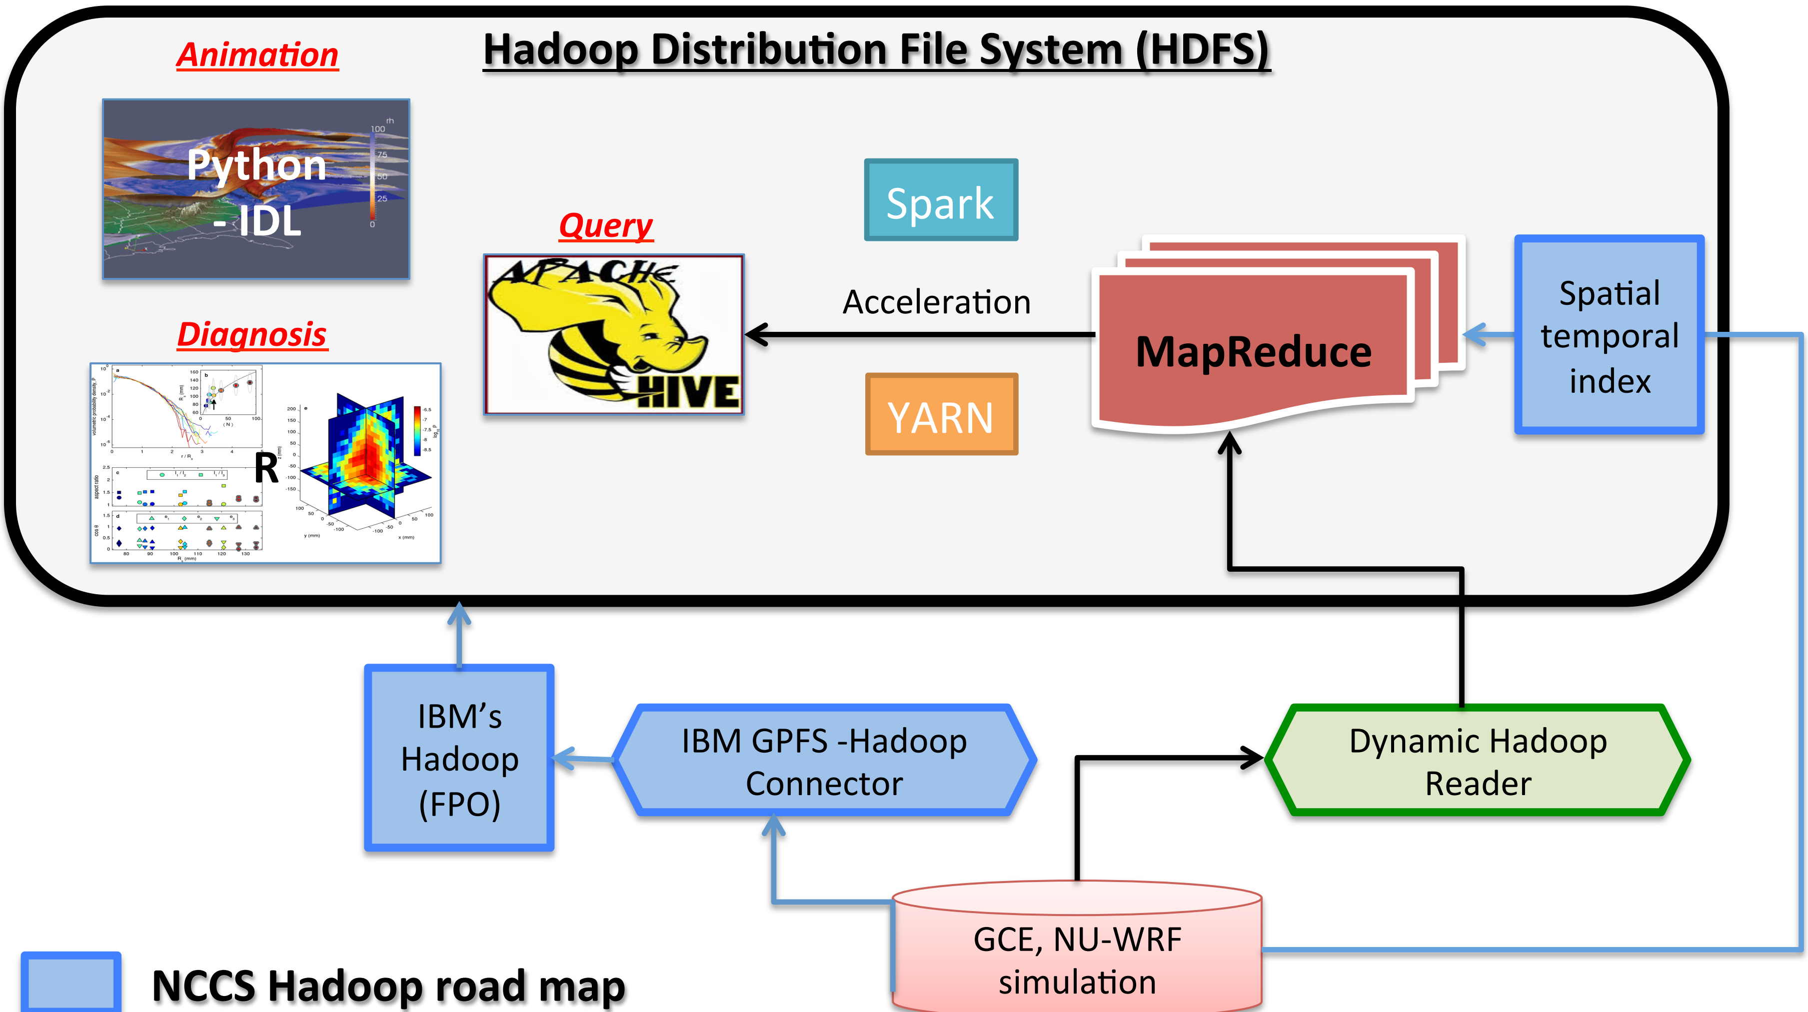

PortHadoop Overview

GPFS (IBM Spectrum Scale 4.2.0.1)
PortHadoop (based on Cloudera Hadoop 5.3.3)

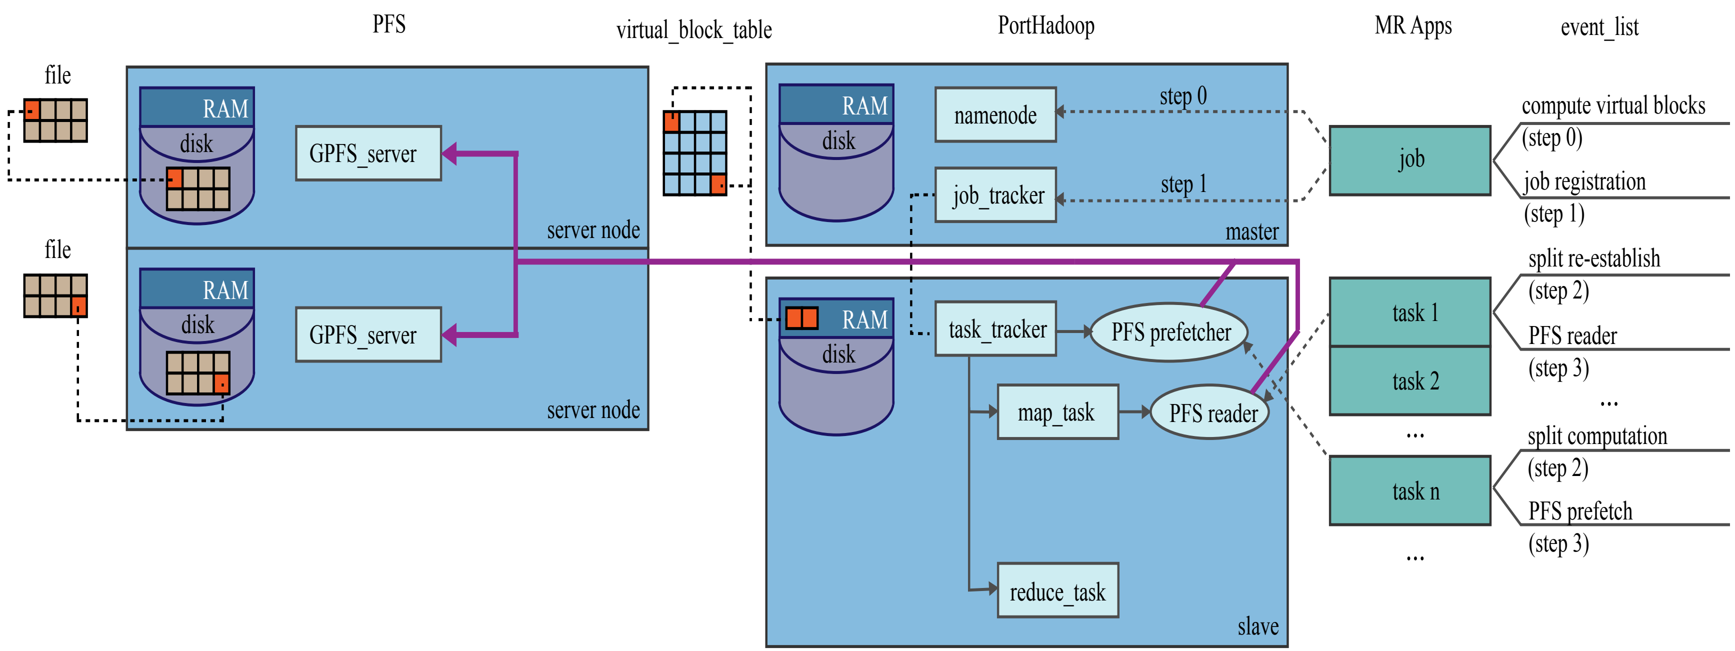

Adpatively Subsetting: Data Flow

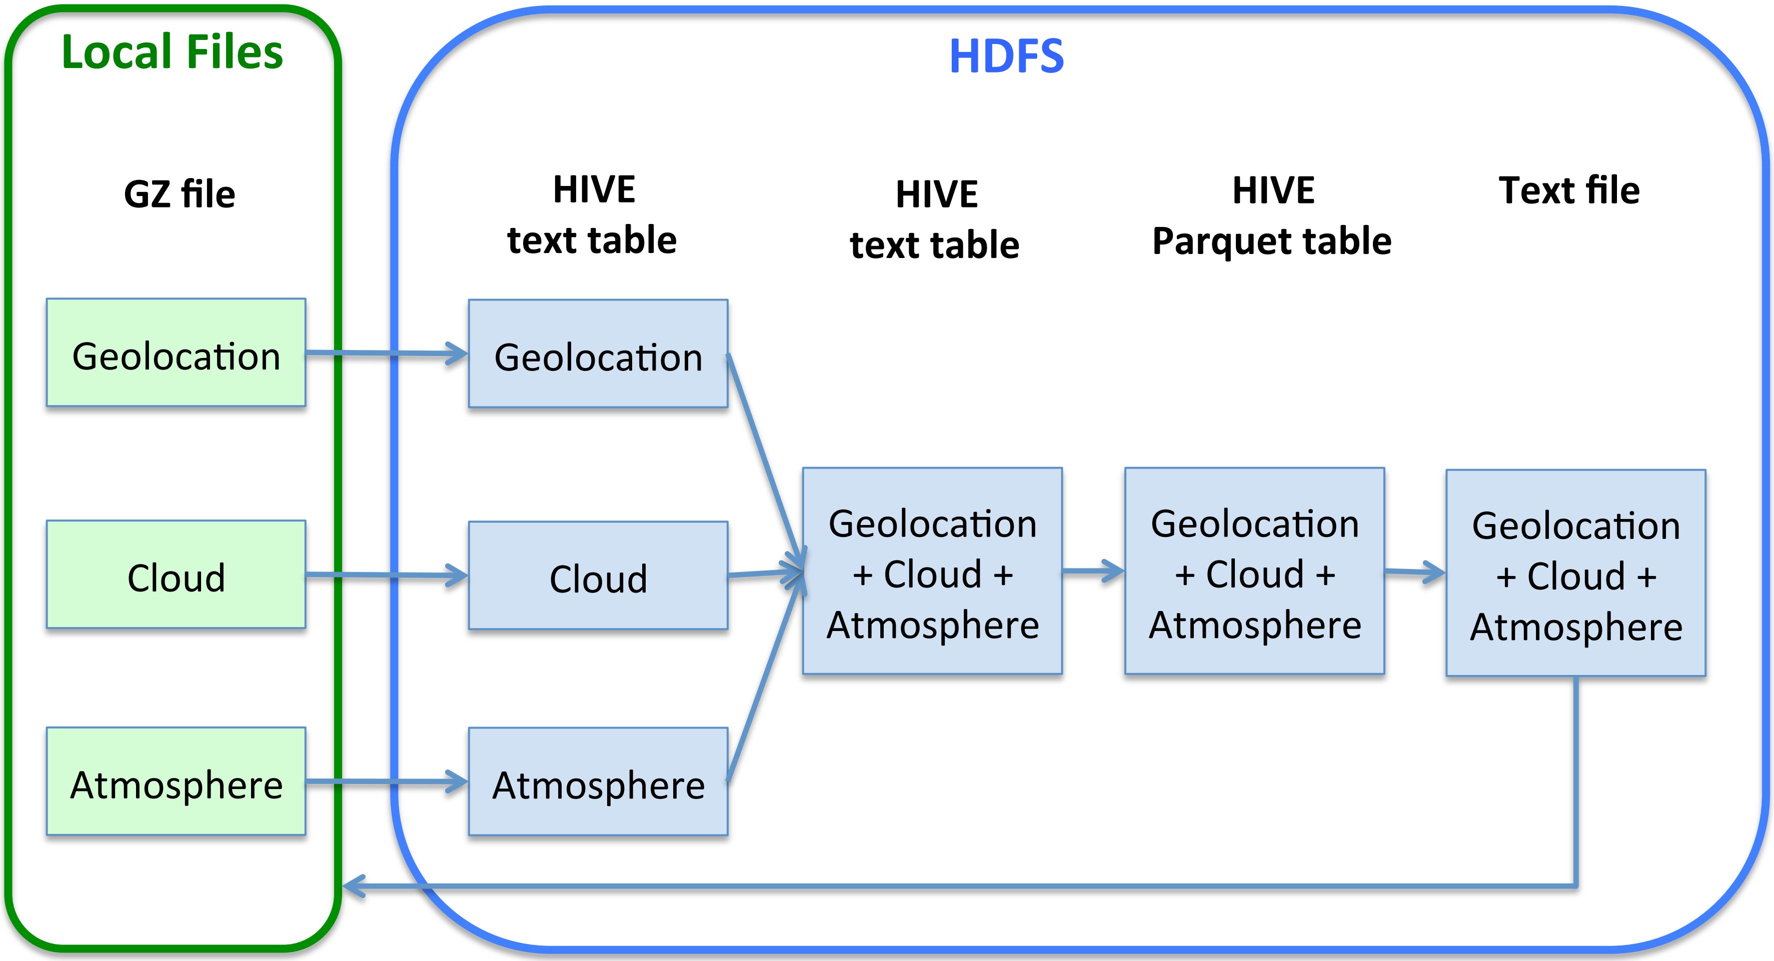

Using multiple tables saves storage as well as be flexible.
Combining tables is compute-intensive.

Visualization and Diagnosis via MapReduce

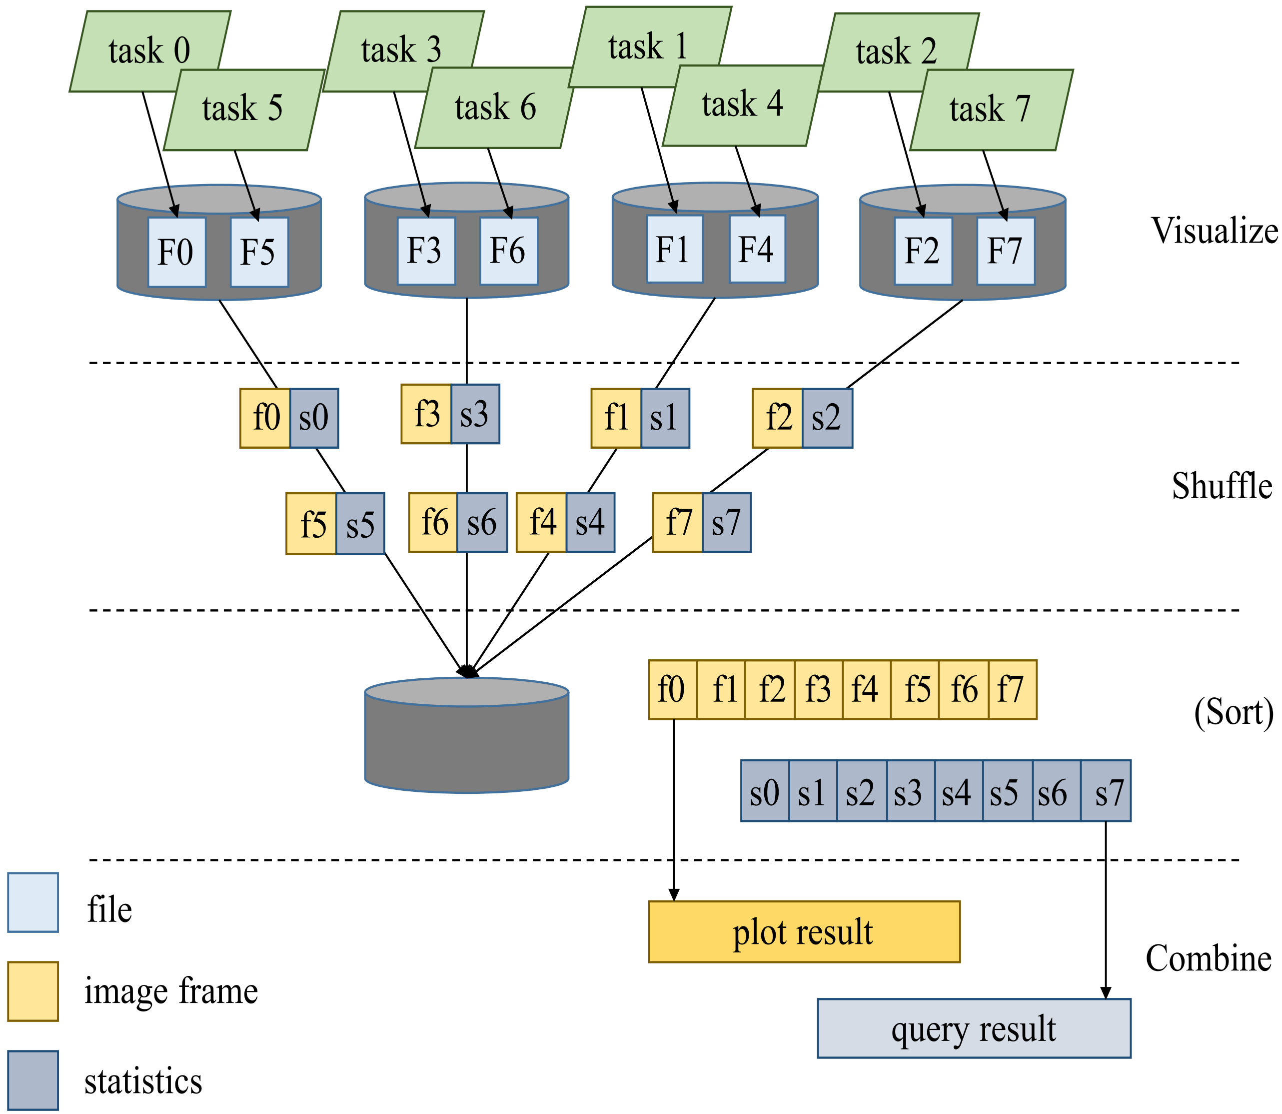

Submitted to The 6th IEEE International Conference on Big Data and Cloud Computing (BDCloud 2016)

Visualization of NU WRF 1250x1250 Rain Simulation woth R+Hadoop and Spark R

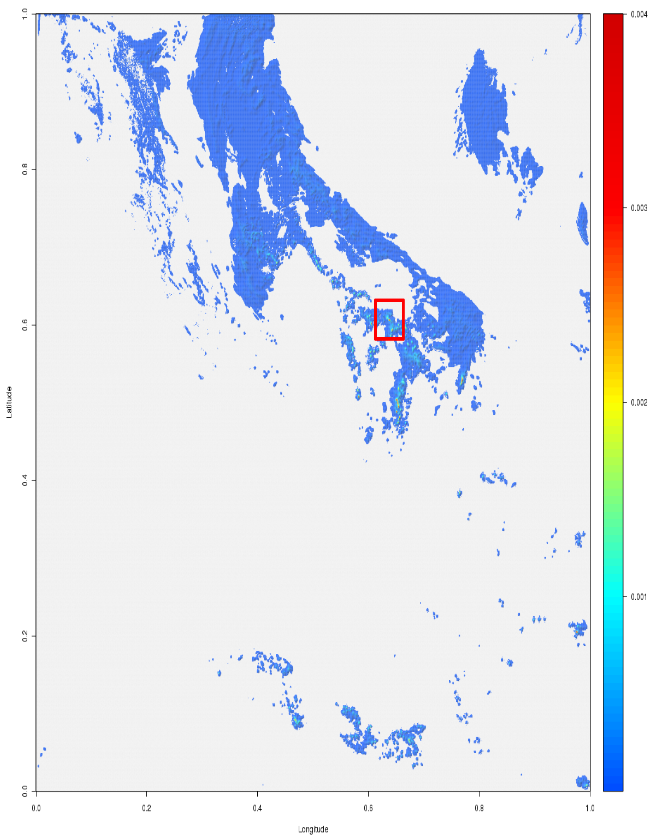

An image with a highlighted area for an interested event (the heaviest rainfall)

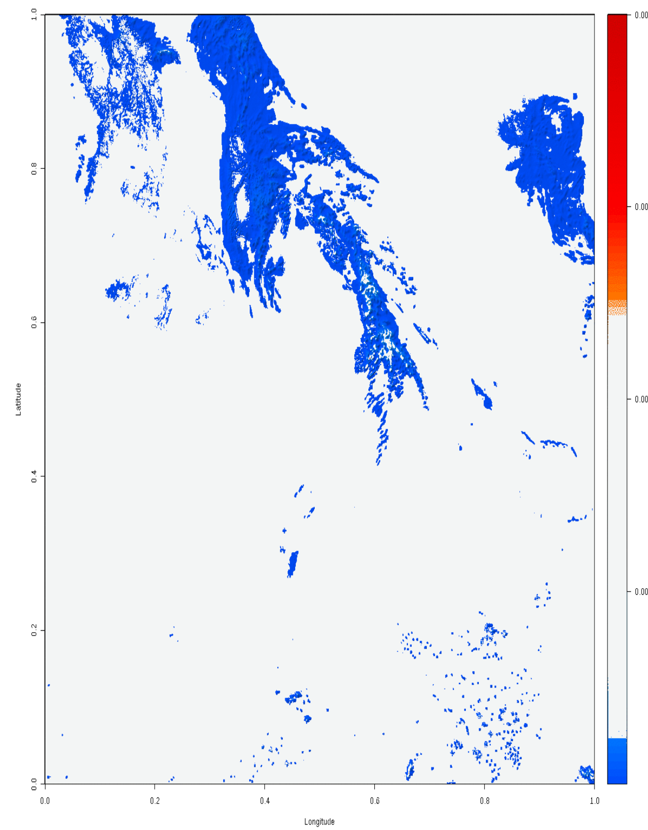

Animation (From 9 AM, 2014-04-28 to 12 AM, 2014-04-29)

Performance with Spark R

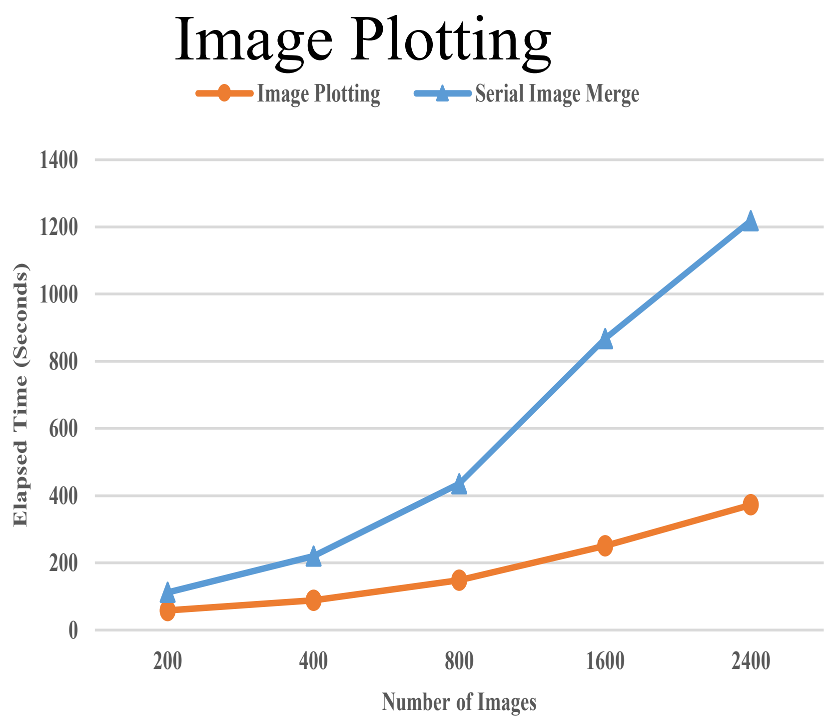

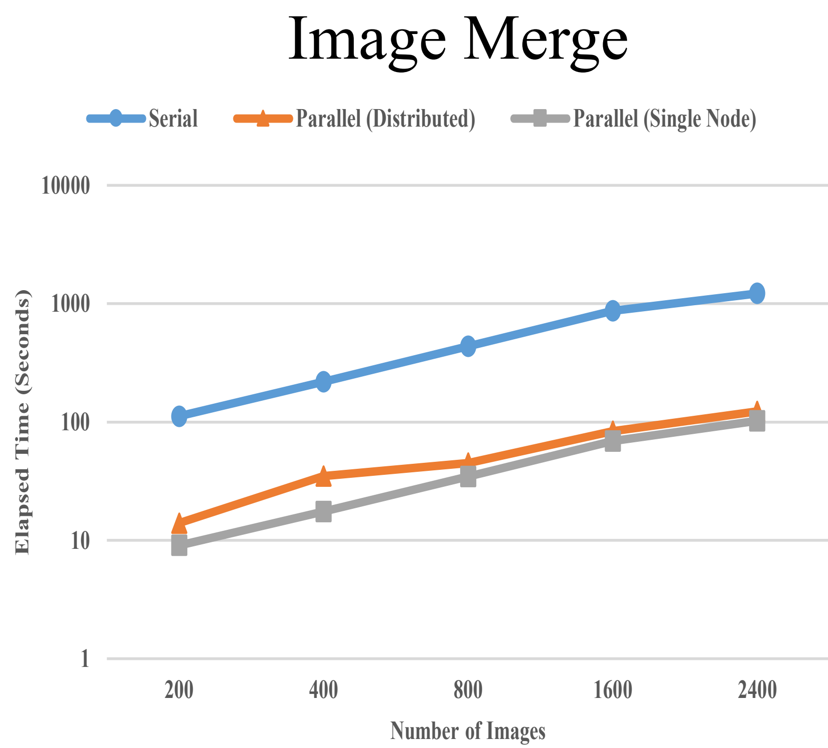

Data: NU-WRF model with a 1250×1250×50 grid (4km resolution) and 48-hour simulation time. Each time frame has ~3GB data. An image is created for each layer

Testbed: 9 nodes (1 master + 8 slaves). Each node has 48 cores and 128 GB memory.

Method: Spark R

Query and Adaptively Subsetting

TABLE II: Simple SQL query statements

Labels in Figure 10	SQL statements
Top 1	select * from dataframe where value == (select max(value) from dataframe) limit 1
Top 10	select * from dataframe desc order by value limit 10
Condition	select * from dataframe where value > 0.005

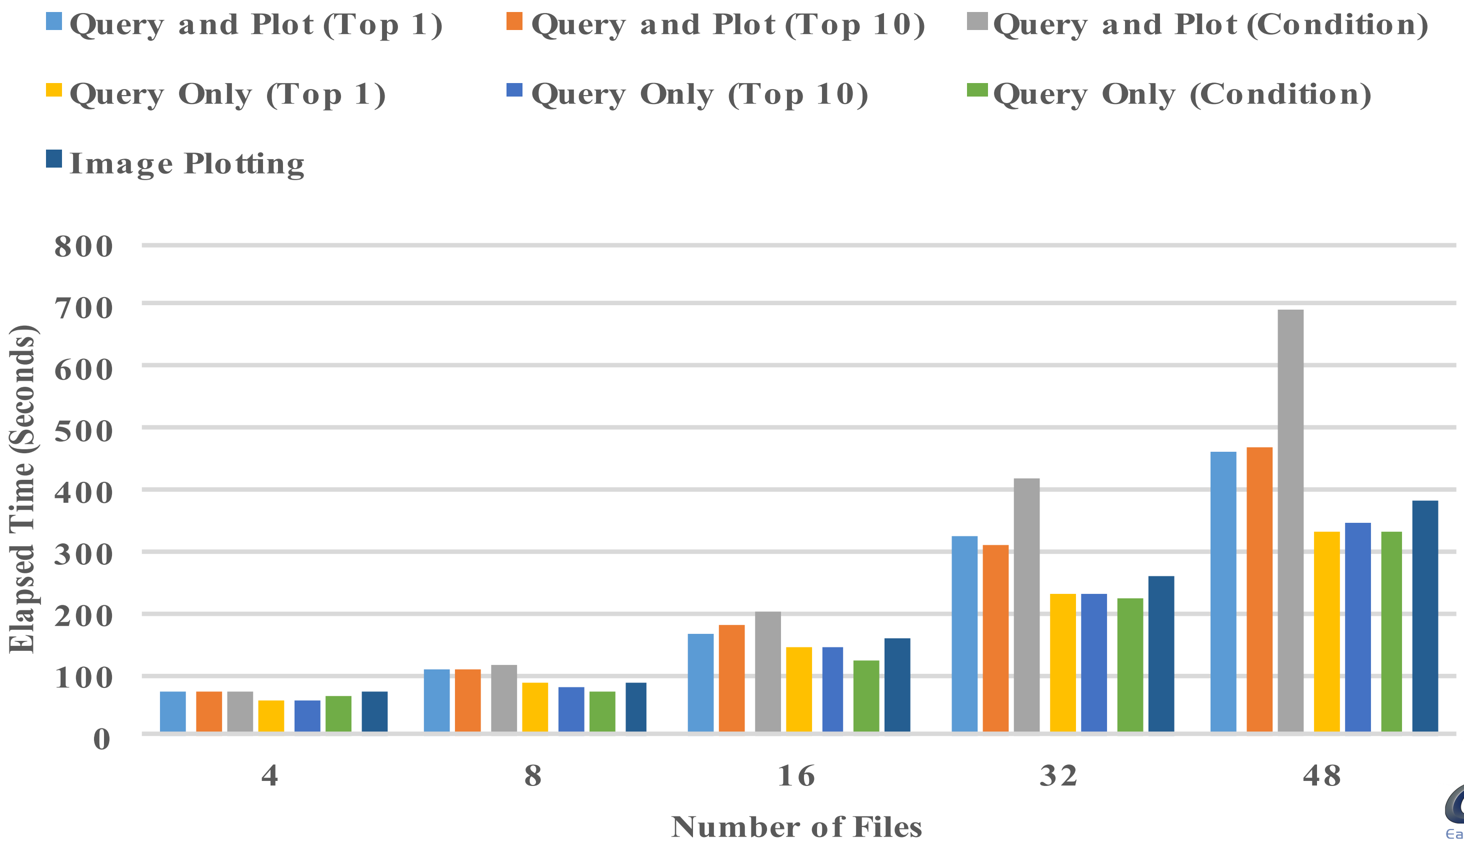

SCL Web Design

SCL Users Defined

SCL User Hierarchy

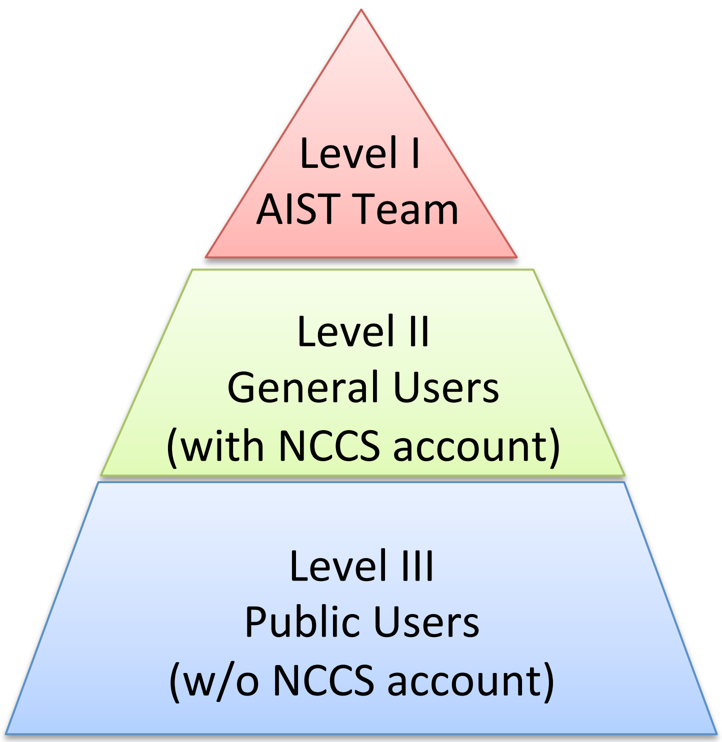

User Level (#)	Upload new data	Visualization	Statistics	Subset	Download data	Add new statistical function
I (~5+)	Yes	Yes	Yes	Yes	Yes	Yes
II (~40)	Yes, upon request to NCCS[1]	Yes	Yes	Yes	Yes	N/A
III (~1000)	N/A	Yes[2]	Yes[2]	Yes[2]	Yes	N/A

[1] after discussion with Dan Duffy@NCCS.

[2] This depends on Hadoop resource and NASA security.

Super Cloud Library Web Design

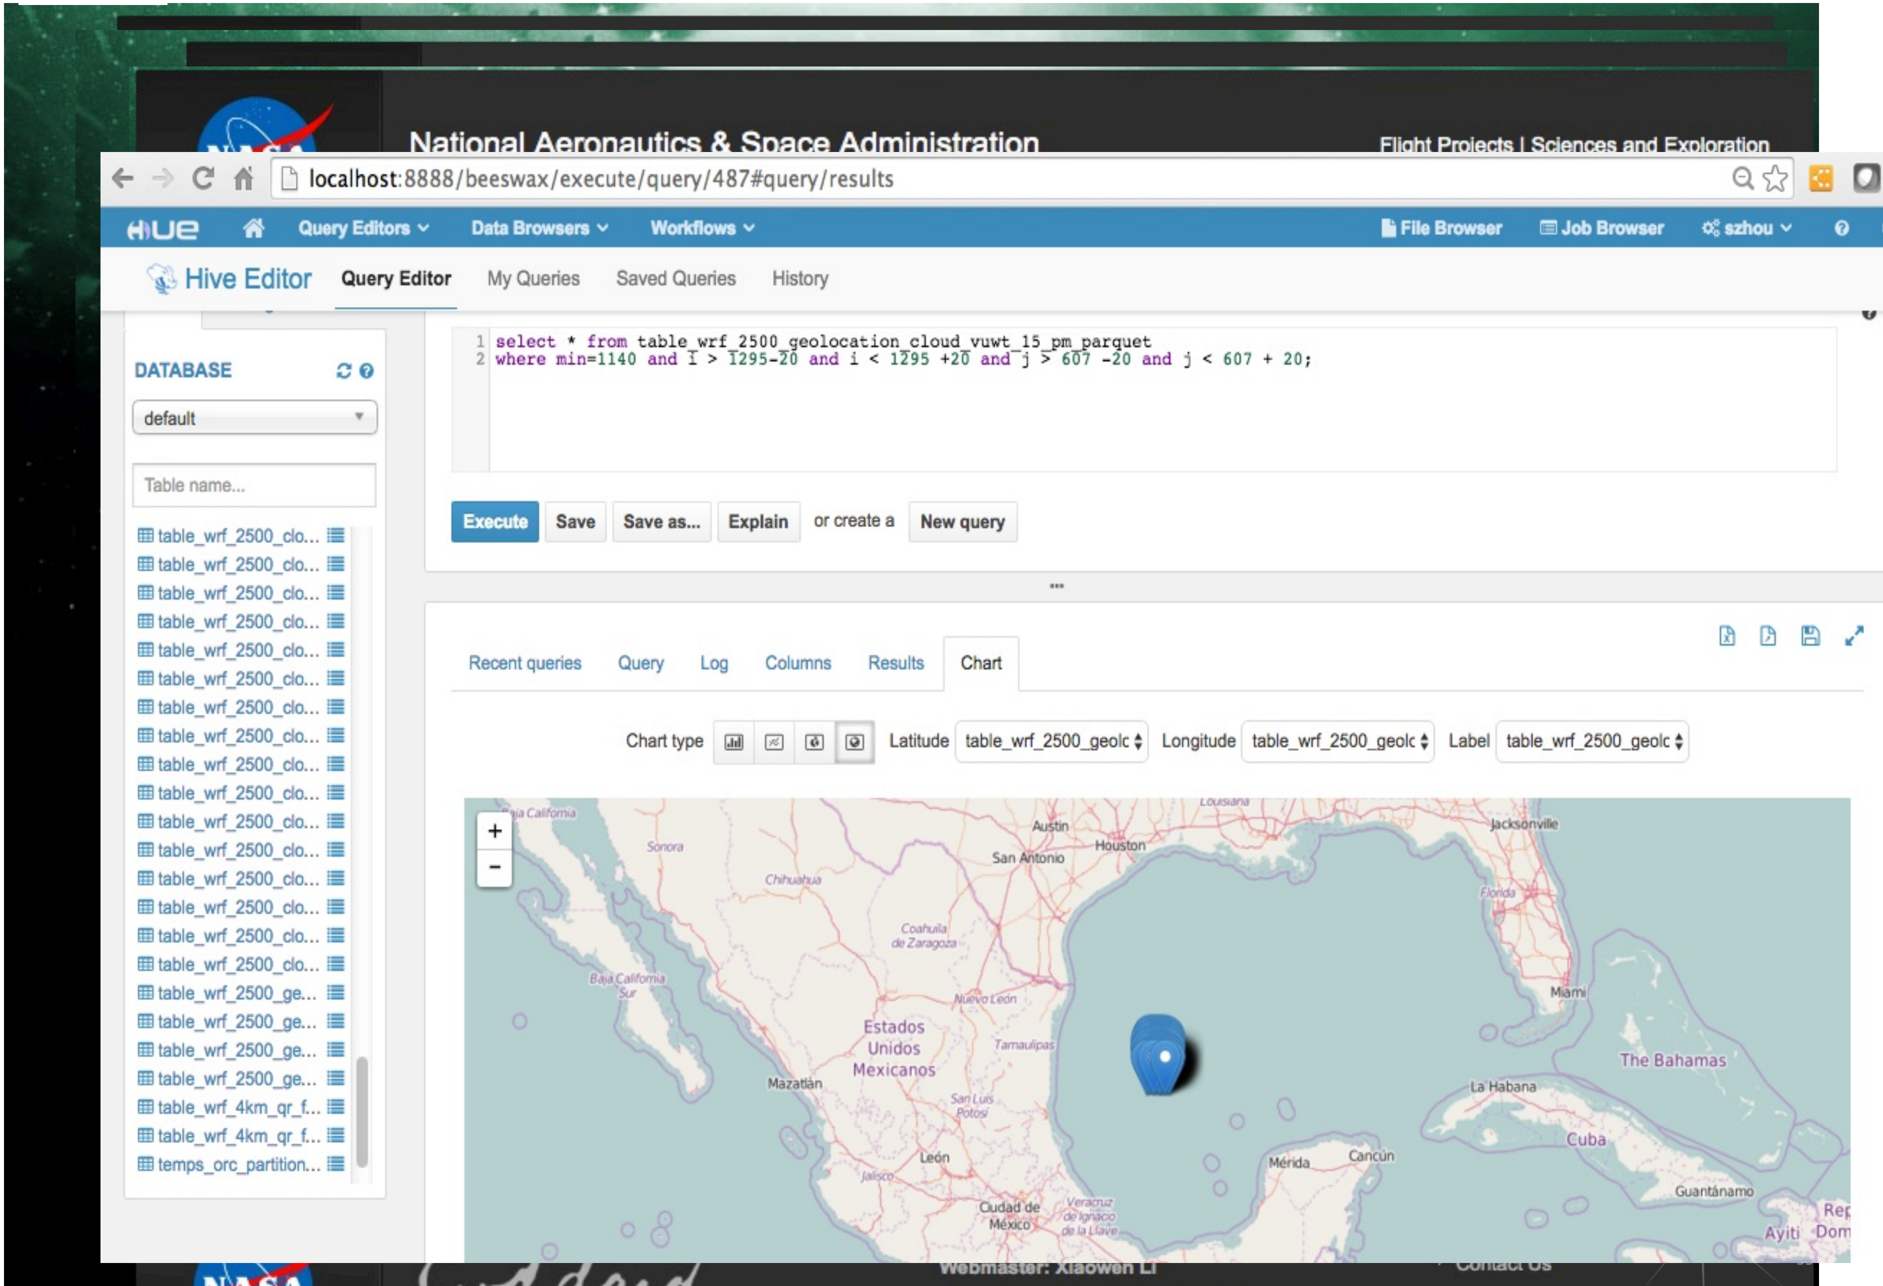

Summary/ Future Work

- Develop Super Cloud Library (SCL) supporting Cloud Resolving Model Data Analyses using Spark on Hadoop.
  - *Create cloud data files : Model inter-comparison*
  - *Develop data model and Hadoop format transformer: Improvement for performance*
  - *Develop a dynamic Hadoop reader tool: NCCS*
  - Develop subset and visualization APIs (Application Programming Interfaces): Tested and need work on diagnosis analyses
  - Develop a Web User Interface: Proposed SCL website
- Conduct **Demo** of GCE and NU-WRF diagnoses on NCCS: By February 2017