



# Introduction to Big Data Analytics using NEXUS

## 2017 ESIP Federation Summer Meeting Workshop

Jet Propulsion Laboratory  
California Institute of Technology

POC: [thomas.huang@jpl.nasa.gov](mailto:thomas.huang@jpl.nasa.gov)

Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not constitute or imply its endorsement by the United States Government or the Jet Propulsion Laboratory, California Institute of Technology.



# IN009: Big Data Analytics

## IN009: Big Data Analytics

Submit an Abstract to this Session

**Session ID#:** 24747

**Session Description:**

Big Data pose great challenges for Earth and Space sciences. Cloud Computing emerged as a promising solution for supporting Big Data analytics in areas such as climate science, ocean science, atmospheric science, planetary science, and other geoscience domains for model simulation, data management, information mining, decision support, knowledge discovery and visualization. As a follow-on to the 2016 success at AGU, this session is to capture the latest on applying Cloud Computing for Big Data Analytical problems in all Earth and space domains. Topics include experiments, demonstration, studies, methods, solutions and solution discussion on:

Solutions for big data analytics

Big data management and mining

Application of open source technologies

Automated techniques for data analysis

Browser-based data analytics and visualization

Real time decision support

Contributions that fuse participatory social learning into the Geoscience R&D processes are also welcome.

**Primary Convener:**

**Thomas Huang**, NASA Jet Propulsion Laboratory, Pasadena, CA, United States

**Conveners:**

**Chaowei Phil Yang**, George Mason University Fairfax, Fairfax, VA, United States, **Tiffany C**

**Vance**, NOAA Seattle, Seattle, WA, United States and **Brian D Wilson**, Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA, United States

**Index Terms:**

1914 Data mining [INFORMATICS]

1916 Data and information discovery [INFORMATICS]

1928 GIS science [INFORMATICS]

1932 High-performance computing [INFORMATICS]



[https://agu.confex.com/agu/fm17/  
preliminaryview.cgi/Session24747](https://agu.confex.com/agu/fm17/preliminaryview.cgi/Session24747)



# Agenda

- Overview
- Applications
- Scale for Speed
- Jupyter Notebook and NEXUS
- Hands-on Labs



# Overview

**Thomas Huang**

Jet Propulsion Laboratory  
California Institute of Technology



## Reality

- With large amount of observational and modeling data, downloading to local machine is becoming inefficient
- Data centers are starting to provide additional services
  - Better searches – faceted, spatial, keyword, relevancy, etc.
  - Data subsetting – data reduction
  - Visualization – visual discovery

## 2015 NASA ESTO/AIST Big Data Study Roadmap: Moving from Data Archiving to Data Analytics

### Increasing “big data” era is driving needs to

- Scale computational and data infrastructures
- Support new methods for deriving scientific inferences
- Shift towards integrated data analytics
- Apply computational and data science across the lifecycle

### Scalable Data Management

- Capturing well-architected and curated data repositories based on well-defined data/information architectures
- Architecting automated pipelines for data capture

### Scalable Data Analytics

- Access and integration of highly distributed, heterogeneous data
- Novel statistical approaches for data integration and fusion
- Computation applied at the data sources
- Algorithms for identifying and extracting interesting features and patterns




# Facts

- Moving / copying science data (and managing copies) is more expensive than computation
- Hardware & software do not yet make science data analysis easy at terabyte scales
- Current analytics are mostly I/O bound
- Next generation – “advanced” analytics will be compute bound (simulation, distributed linear algebra)
- Current files formats are good for data archival, NOT for data analysis

**"The scientific file-formats of HDF, NetCDF, and FITS can represent tabular data but they provide minimal tools for searching and analyzing tabular data... Performing this filter-then-analyze, data analysis on large datasets with conventional procedural tools runs slower and slower as data volumes increase."**

**-- Jim Gray, Scientific Data Management in the Coming Decade**

# Traditional Data Analysis



- Depending on the data volume (size and number of files)
- It could take many hours of download – (e.g. 10yr of observational data could yield thousands of files)
- It could take many hours of computation
- It requires expensive local computing resource (CPU + RAM + Storage)
- After result is produced, purge downloaded files

## Observation

- Traditional methods for data analysis (time-series, distribution, climatology generation) can't scale to handle large volume, high-resolution data. They perform poorly
- Performance suffers when involve large files and/or large collection of files
- A high-performance data analysis solution must be free from file I/O bottleneck


# NASA's Upcoming Big Data Mission: Surface Water and Ocean Topography (SWOT)

**Oceanography:** Characterize the ocean mesoscale and sub-mesoscale circulation at spatial resolutions of 10 km and greater.

**Hydrology:** To provide a global inventory of all terrestrial water bodies whose surface area exceeds  $(250\text{m})^2$  (lakes, reservoirs, wetlands) and rivers whose width exceeds 100 m (requirement) (50 m goal) (rivers).

- To measure the global storage change in fresh water bodies at sub-monthly, seasonal, and annual time scales.
- To estimate the global change in river discharge at sub-monthly, seasonal, and annual time scales.

- **Data Volume:**
  - 17PB of original data
  - 6 PB of reprocessed data
- **Total of about 23PB for a nominal 3-year mission**
- **Add roughly 450TB/month for any mission extension**



Launches April of 2021  
<https://swot.jpl.nasa.gov>


# The Silver Bullet?



- Provides a elastic infrastructural approach to Big Data
- Convenient, rapid access to a shared pool of computing resources
- It can be strictly infrastructure (networks, servers, storage, etc.) or can be combined with software to facilitate ready access to applications and services
- It is not just resource virtualization with a new name; features such as self-service provisioning and advanced metering of use set it apart as a new and transformational technology

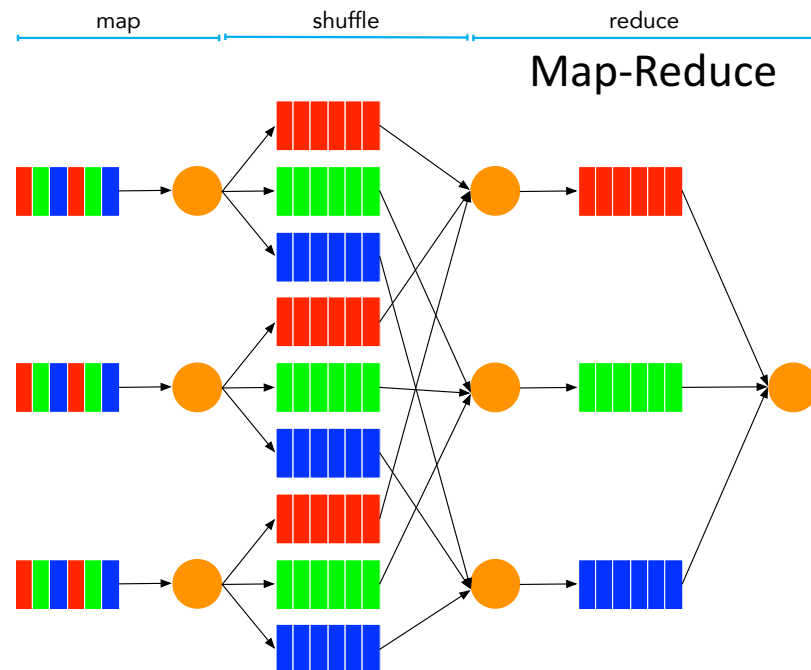


# But... How Do We Get There?



# Solution that Scales

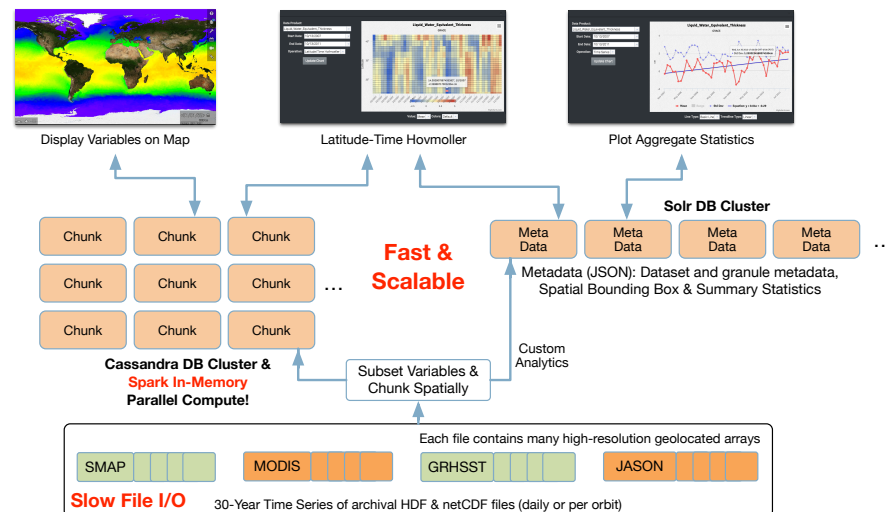
- From the ground up, NEXUS was designed to take advantage of multiple computing node and cloud-based storage
- Ingestion**
  - Supports multiple processing streams
  - Add new streams without reboot
  - Support various datasets
- Data Management**
  - Cluster-based geospatial index search
  - Cluster-based NoSQL storage
  - Object-store backend
- Analytics**
  - Apache Spark analytic platform for auto scaling
  - Supports multi-processing for algorithms that are not suitable for map-reduce



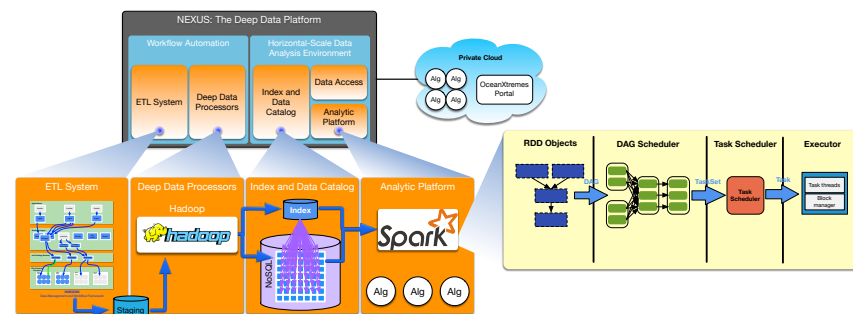
- Map-Reduce provides a simple, manageable framework to scale-out processing
- However, it requires data to be partitioned in small blocks
- A typical block in Hadoop is 64MB to 128MB
- A single granule could be hundred of megabytes to several gigabytes
- To take advantage of the Map-Reduce framework, we need to divide granule into blocks

# NEXUS: Deep Data Platform

- A data-intensive analysis solution using a new approach for handling science data to enable large-scale data analysis
- Streaming architecture for horizontal scale data ingestion
- Scales horizontally to handle massive amount of data in parallel
- Provides high-performance geospatial and indexed search solution
- Provides tiled data storage architecture to eliminate file I/O overhead
- A growing collection of science analysis webservices using Apache Spark: parallel compute, in-memory map-reduce framework
- Pre-Chunk and Summarize Key Variables
  - Easy statistics instantly (milliseconds)
  - Harder statistics on-demand using Spark (in seconds)
  - Visualize original data (layers) on a map quickly (Cassandra store)
- **Algorithms** – Time Series | Latitude/Time Hovmöller| Longitude/Time Hovmöller| Latitude/Longitude Time Average | Area Averaged Time Series | Time Averaged Map | Climatological Map | Correlation Map | Daily Difference Average



Two-Database Architecture

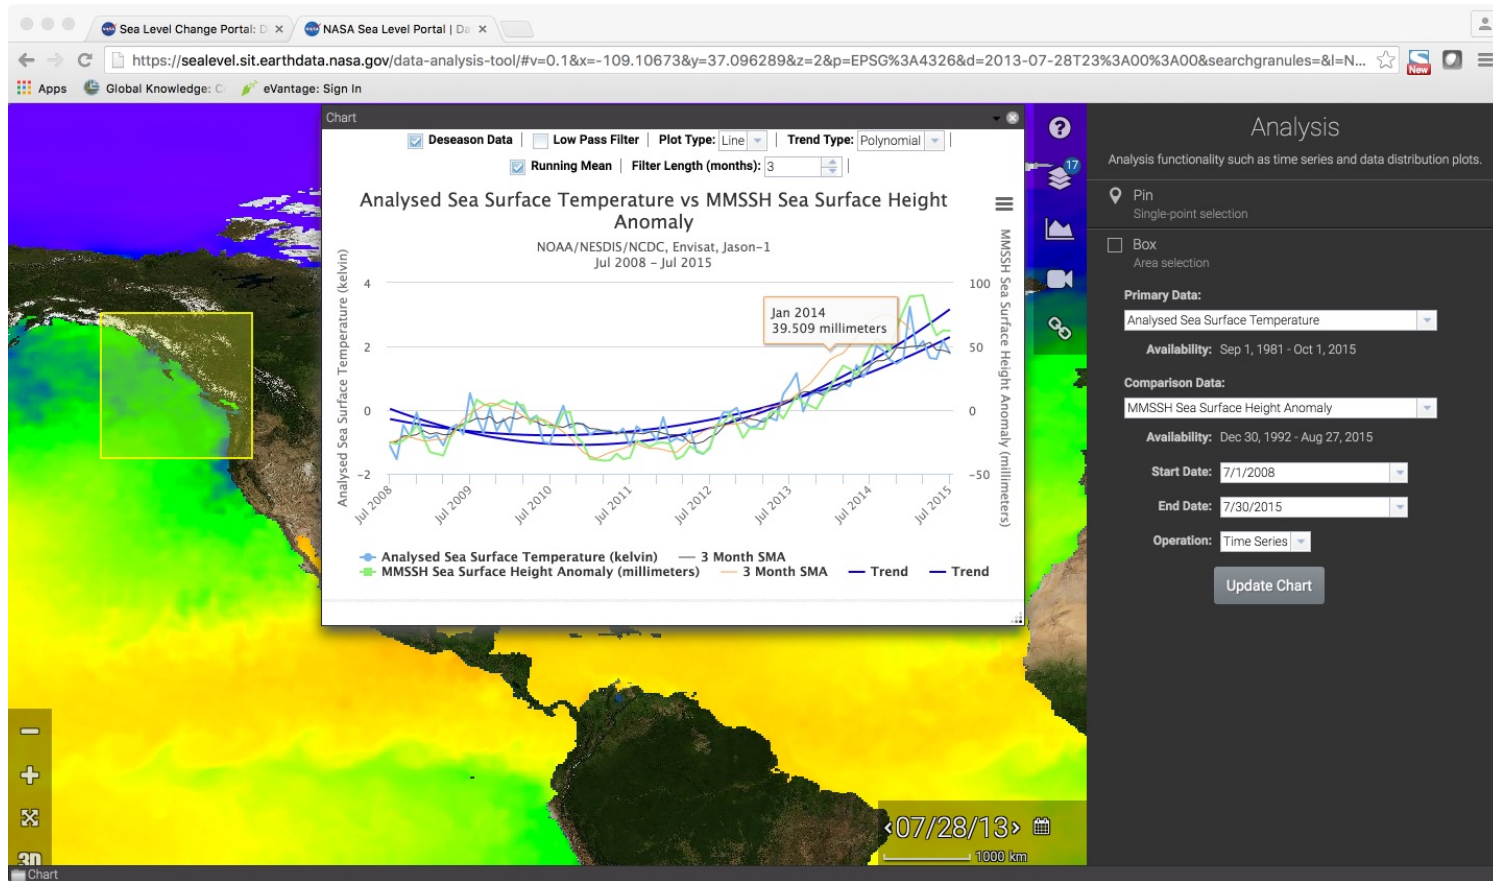


Deep Data Computing Environment (DDCE)

**Open Source: Apache License 2**

<https://github.com/dataplumber/nexus>





"The Blob is a result of a high pressure system that has parked itself in the Gulf of Alaska for the past few years that has driven the polar jet stream north into northern Canada and then it plunged rapidly out of northern Canada into the American Midwest and northeast. And so the result was hot dry winters on the west coast, and fierce winters with heavy snow pack in the Midwest." – Bill Patzert, NASA/JPL



# Jupyter Notebook Integration

<https://jupyter.jpl.nasa.gov>

```
# Request NEXUS to compute SST Time Series 2008/9/1 - 2015/10/1
# for the "blob" warming off Western Canada and plot the means
...
ds='AVHRR_OI_L4_GHRSSST_NCEI'

url = ... # construct the webservice URL request


# make request to NEXUS using URL request
# save JSON response in local variable
ts = json.loads(str(requests.get(url).text))

# extract dates and means from the response
means = []
dates = []
for data in ts['data']:
    means.append (data[0]['mean'])
    d = datetime.datetime.fromtimestamp((data[0]['time']))
    dates.append (d)

# plot the result
...
```

[https://oceanxtremes.jpl.nasa.gov/timeSeriesSpark?spark=measos,16,32&ds=AVHRR\\_OI\\_L4\\_GHRSSST\\_NCEI&minLat=45&minLon=-150&maxLat=60&maxLon=-120&startTime=1220227200&endTime=1443657600](https://oceanxtremes.jpl.nasa.gov/timeSeriesSpark?spark=measos,16,32&ds=AVHRR_OI_L4_GHRSSST_NCEI&minLat=45&minLon=-150&maxLat=60&maxLon=-120&startTime=1220227200&endTime=1443657600)

It took: 6.909193992614746 sec

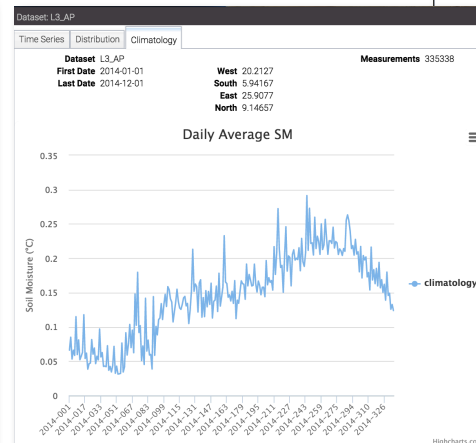
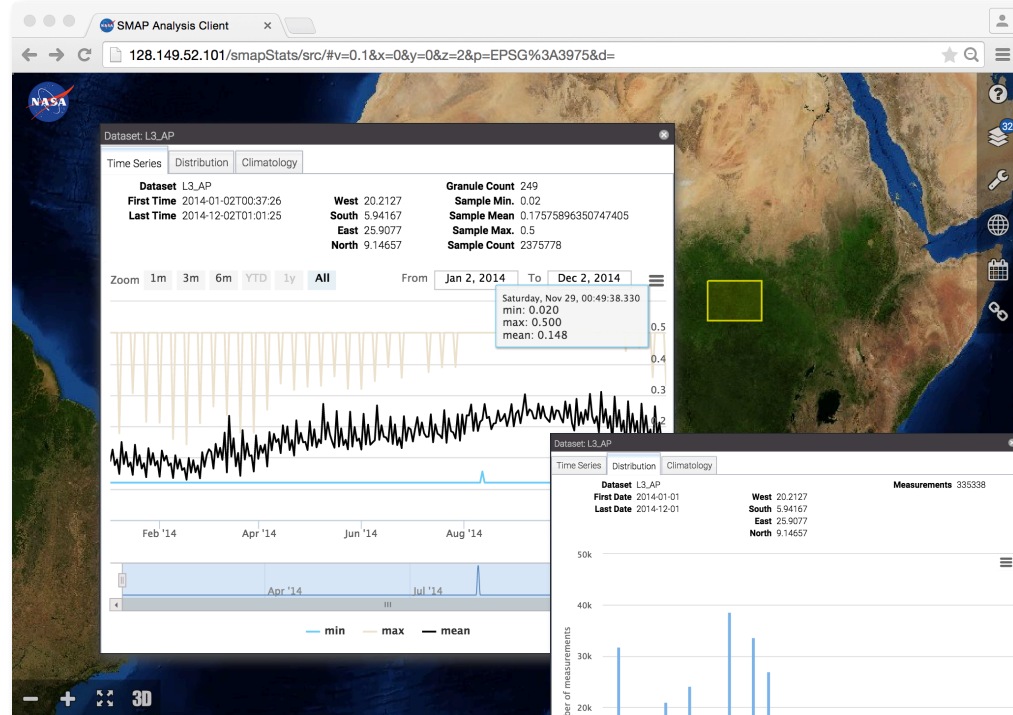
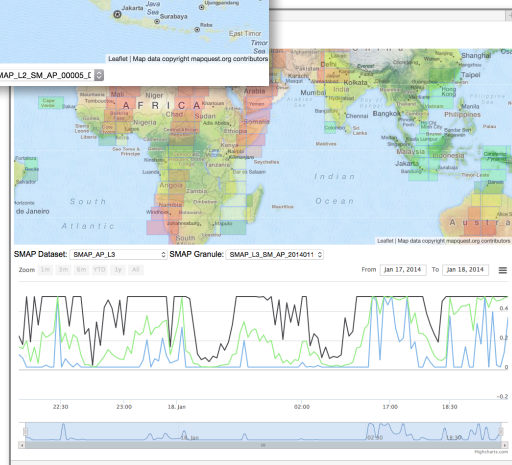
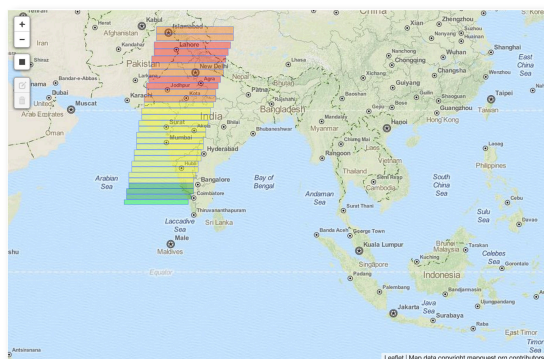




# Applications of NEXUS

# SMAP Cal/Val Data Analysis

- Application of NEXUS data tiling architecture for analyzing SMAP L2 and L3 data
- Provide on-the-fly analysis of L2 and L3 datasets
- Provides on-the-fly data subsetting using data tiling architecture

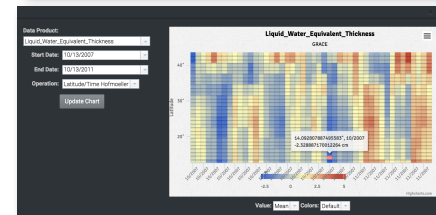

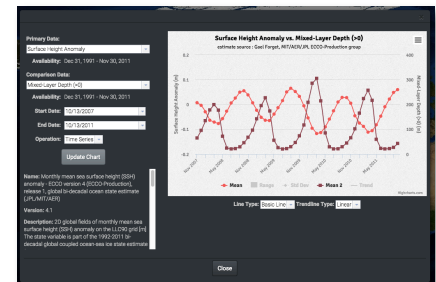
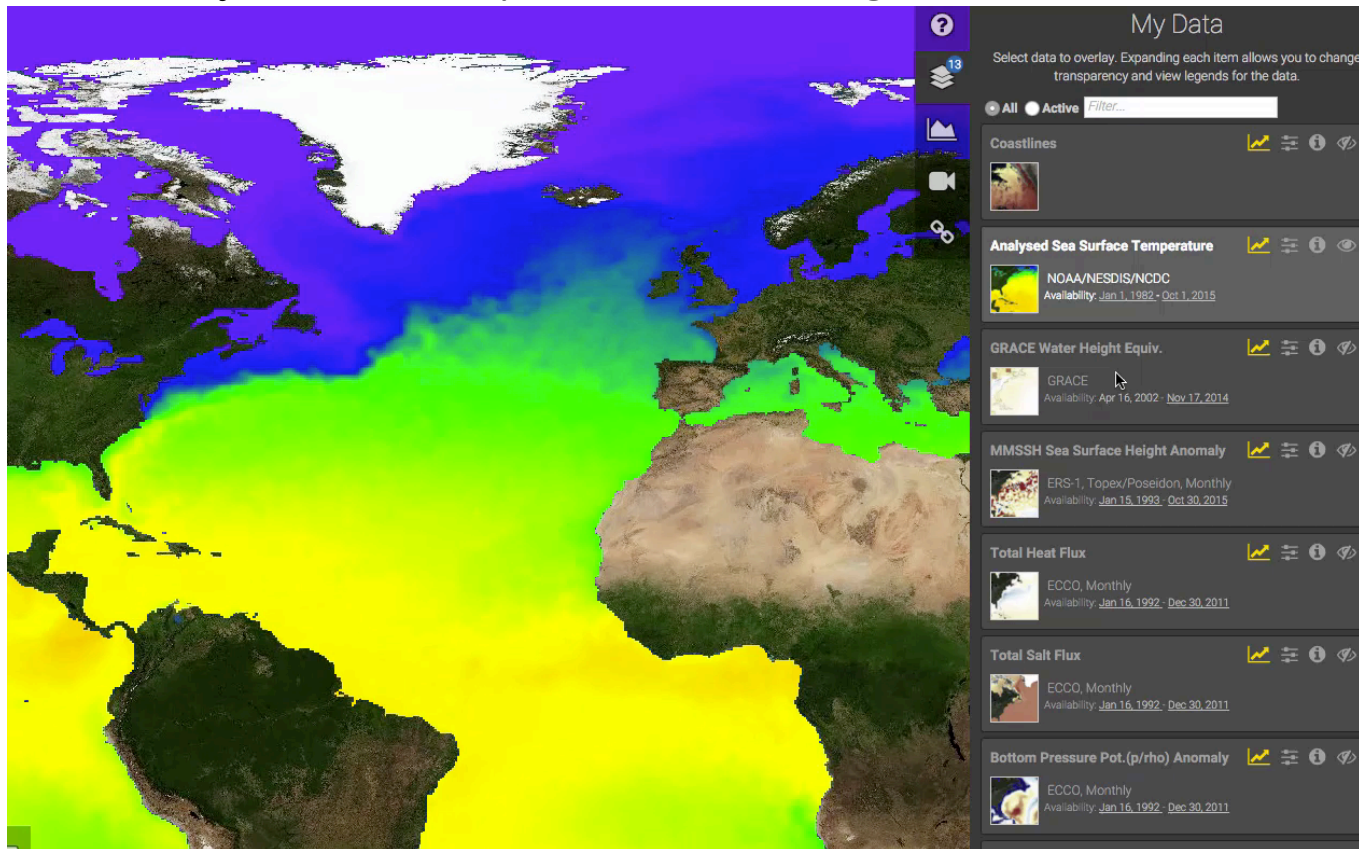






# SEA LEVEL CHANGE Observations from Space

Data Analysis Tool <https://sealevel.nasa.gov>



Facebook: 28K followers

Twitter: 22K followers

**“NASA Sea Level Change Website Offers Everything You Need To Know About Climate Change”**

<http://www.techtimes.com/articles/147210/20160405/nasa-sea-level-change-website-offers-everything-need-know-climate.htm>

**“NASA’s new sea level site puts climate change papers, data, and tools online”**

<http://techcrunch.com/2016/04/04/nasas-new-sea-level-site-puts-climate-change-papers-data-and-tools-online/>

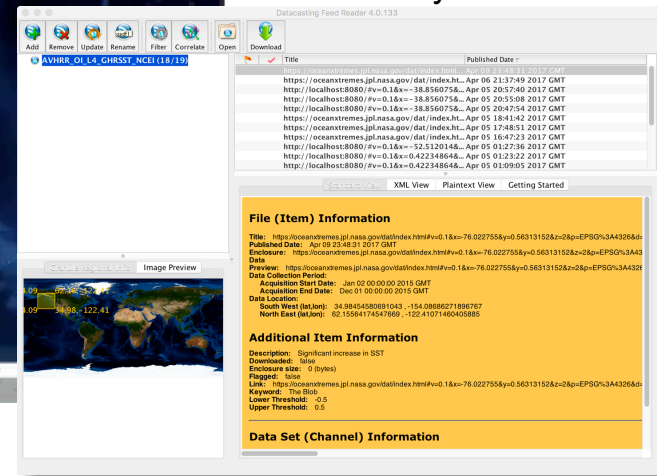


# OceanXtremes: Oceanographic Data-Intensive Anomaly Detection and Analysis Portal



<https://oceanxtremes.jpl.nasa.gov>

- An oceanographic data-intensive anomaly detection and analysis portal
- Cloud-based big data analytic platform for
  - Climatology generation
  - On-the-fly daily difference computation
  - Anomaly registry and publication
  - On-the-fly data analytics



DataCasting client is able to pickup Anomaly Cast published by OceanXtremes

# Hurricane Katrina Study – Using OceanXtremes


Hurricane Katrina passed to the southwest of Florida on Aug 27, 2005. The ocean response in a 1 x 1 deg region is captured by a number of satellites. The initial ocean response was an immediate cooling of the surface waters by 2 deg C that lingers for several days. Following this was a short intense ocean chlorophyll bloom a few days later. The ocean may have been “preconditioned” by a cool core eddy and low sea surface height.

The SST drop is correlated to both wind and precipitation data. The Chl-A data is lagged by about 3 days to the other observations like SST, wind and precipitation.

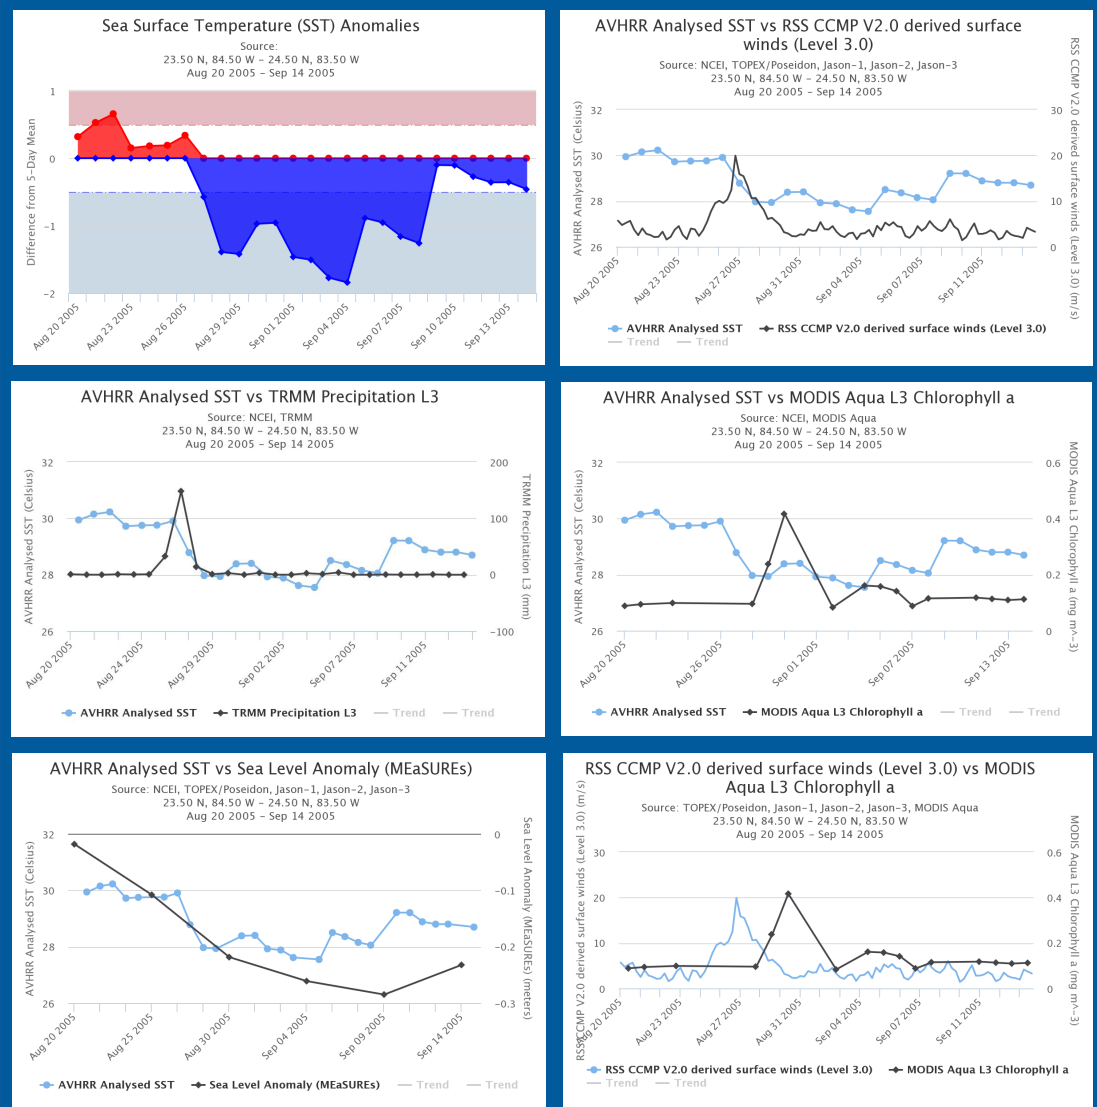
*A study of a Hurricane Katrina-induced phytoplankton bloom using satellite observations and model simulations*  
Xiaoming Liu, Menghua Wang, and Wei Shi

JOURNAL OF GEOPHYSICAL RESEARCH, VOL. 114,  
C03023, doi:10.1029/2008JC004934, 2009

[http://shoni2.princeton.edu/ftp/lyo/journals/Ocean/phybiog\\_eochem/Liu-et-al-KatrinaChlBloom-JGR2009.pdf](http://shoni2.princeton.edu/ftp/lyo/journals/Ocean/phybiog_eochem/Liu-et-al-KatrinaChlBloom-JGR2009.pdf)



Hurricane Katrina  
TRMM overlay SST Anomaly



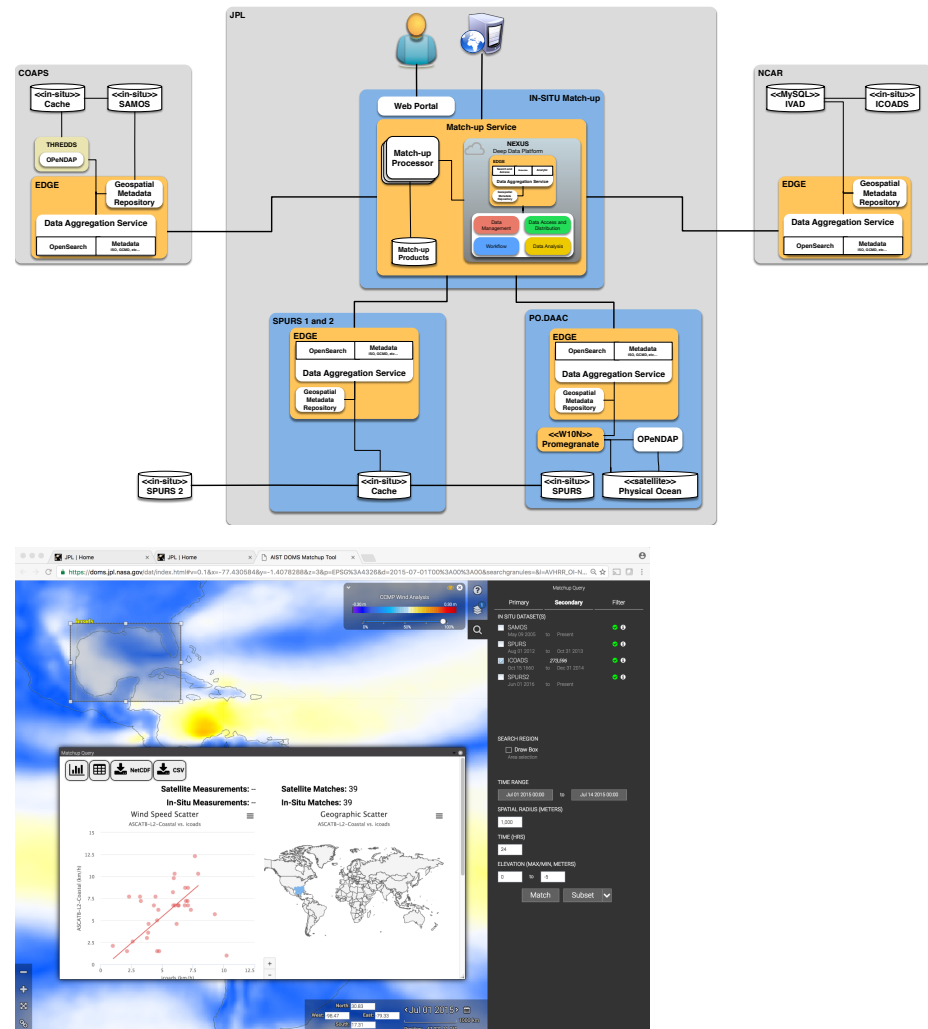
Powered By NEXUS





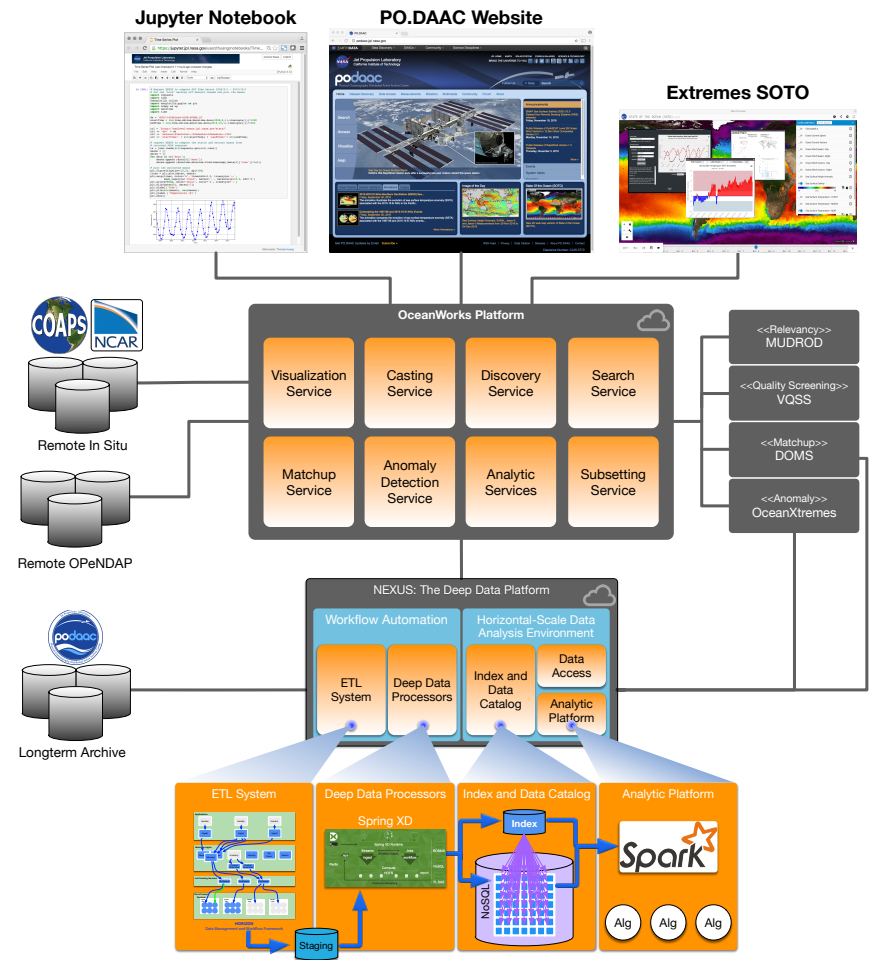

# DOMS: Distributed Oceanographic Matchup Service

- <https://doms.jpl.nasa.gov>
- Distributed Oceanographic Matchup Service
- Typically data matching is done using one-off programs developed at multiple institutions
- A primary advantage of DOMS is the reduction in duplicate development and man hours required to match satellite/in situ data
  - Removes the need for satellite and in situ data to be collocated on a single server
  - Systematically recreate matchups if either in situ or satellite products are re-processed (new versions), i.e., matchup archives are always up-to-date.
- In situ data nodes at JPL, NCAR, and FSU operational.
- Provides data querying, subset creation, match-up services, and file delivery operational.
- Prototype graphical user interface (UI) and APIs accessible for external users.
- Plugin architecture for in situ data source using EDGE
  - Extensible Data Gateway Environment is an Apache License 2 open source technology
  - <https://github.com/dataplumber/edge>
- Defined specification for packaging matchup results. Working with Unidata and ESDSWG's data interoperability and standard groups





- **OceanWorks** is to establish an integrated data analytic center at the NASA PO.DAAC for Big Ocean Science. It focuses on technology integration, advancement and maturity
- **Collaboration between JPL, FSU, NCAR, and GMU**
- Bringing together PO.DAAC-related big data technologies
  - **AIST-14 OceanXtremes (PI: Huang/JPL) – TRL 4**  
Anomaly detection and ocean science
  - **NEXUS (PI: Huang/JPL) – TRL 6**  
Deep data analytic platform
  - **AIST-14 DOMS (PI: Smith/FSU) – TRL 4**  
Distributed in-situ to satellite matchup
  - **AIST-14 MUDROD (PI: Yang/GMU) – TRL 7**  
Search relevancy and discovery
  - **ACCESS-13 VQSS (PI: Armstrong/JPL) – TRL 7**  
Virtualized Quality Screening Service



Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not constitute or imply its endorsement by the United States Government or the Jet Propulsion Laboratory, California Institute of Technology.



# NEXUS Upcoming Capabilities

- **NEXUS client library** – wrapper for the RESUTful webservice calls. The library can be integrated into Jupyter notebook or as a standard python module
  - Integration with EOSDIS Cloud Analysis Toolkit to Enable Earth Science (CATEES)
- **Support ESDS-RFC-028v1.1** - Dataset Interoperability Recommendations for Earth Science
- **Support NASA General Application Platform (NGAP)**
- **Object store** – offer object store support for tile storage for cost reduction and infinite storage. Storage change should be transparent to NEXUS users. We do expect small performance reduce compare to NoSQL.
- **Amazon Lambda** – develop new serverless architecture for data ingestion and tiling
- **Amazon EMR integration** – for AWS provision map-reduce solution



# Scale for Speed

**Joseph Jacob**

Jet Propulsion Laboratory  
California Institute of Technology



# Two Data Analysis Platforms at NASA

- **Giovanni**
  - A web service at the NASA DAACs to perform a set of common analytics on NASA data sets
  - File-based input – Reads attributes and data from NetCDF granules
  - Server side implementation based on NetCDF Operators (NCO) Toolkit
    - Optimized compiled C/C++ code
  - Performance is adequate for small datasets, but we are in the era of Big Data
  - Scientific file formats like NetCDF and HDF are great for data archive, but a poor solution for data analysis due to being ill-suited for spatiotemporal search and subsetting.
- **NEXUS**
  - Data analysis platform built from the ground up for big data analytics in the cloud
    - Solr for metadata
    - Cassandra for measurement data
    - Spark for parallel analytics
    - The engine behind several NASA applications
      - Sea Level Change Portal (SLCP)
      - Distributed Oceanographic Match-up Service (DOMS)
      - OceanXtremes – Oceanographic anomaly detection





# Nexus Prototype (ESDIS and JPL)

- NASA Earth Science Data and Information System (ESDIS) sponsored JPL to deploy NEXUS on several cloud platforms and compare performance with Giovanni
  - Goal: Use NEXUS to achieve at least 2X speedup over Giovanni for a set of benchmark algorithms.
- **Benchmark Algorithms:**
  - **Area-Averaged Time Series**  
*Compute the spatial average over a user-specified spatial area for each time step in a user-specified time range.*
  - **Time Averaged Map**  
*Compute an average at each grid coordinate within a user-specified spatial area, averaged over a user-specified time range.*
  - **Correlation Map**  
*Compute the correlation coefficient between two variables over time at each grid coordinate, using simple linear regression.*
  - **User Defined Climatology Map**  
*Compute an average at each grid coordinate for a specified month over specified years.*
- Spatial Scale:
  - Entire globe (or large subset of the globe)
  - State (Colorado)
    - Spans 4 deg in latitude and 7 deg in longitude
  - City (Boulder)
    - Spans 1 deg in both latitude and longitude

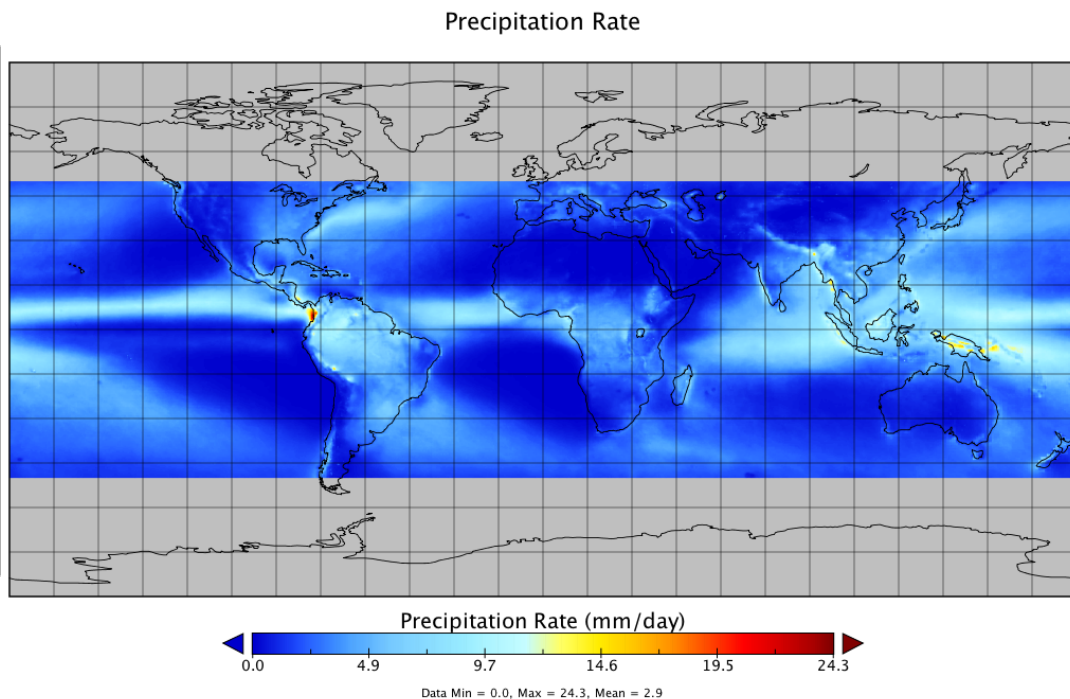

- **Tropical Rainfall Measuring Mission (TRMM)**
  - **Spatial resolution:** 0.25 deg
  - **Spatial coverage:** Entire globe between -50 deg S and +50 deg N latitude
  - **Variable:** TRMM daily precipitation rate (**TRMM\_3B42\_daily\_precipitation\_V7**)
    - Number of daily data granules: 6,574
    - Data volume: 26 GB
    - Dates covered: 1/1/1998 – 12/31/2015
  - **Variable:** TRMM real-time daily precipitation rate (**TRMM\_3B42RT\_daily\_precipitation\_V7**)
    - Number of daily data granules: 5,784
    - Data volume: 23 GB
    - Dates covered: 3/1/2000 – 12/31/2015
  
- **Moderate Resolution Imaging Spectroradiometer (MODIS)**
  - **Spatial resolution:** 1 deg
  - **Spatial coverage:** Entire globe
  - **Variable:** MODIS-Terra Aerosol Optical Depth (AOD) 550 nm dark target (**MOD08\_D3v6**)
    - Number of daily data granules: 5,789
    - Data volume: ~3 GB
    - Dates covered: 3/1/2000 – 2/29/2016
  - **Variable:** MODIS-Aqua Aerosol Optical Depth (AOD) 550 nm dark target (**MYD08\_D3v6**)
    - Number of daily data granules: 5,789
    - Data volume: ~3 GB
    - Dates covered: 3/1/2000 – 2/29/2016

# Time Series: Comparison of Giovanni and NEXUS Performance



- NEXUS run on 8-node cluster computer at JPL running Solr, Cassandra, Spark 2.0, Mesos
- Area-Averaged Time Series over the continental United States
- Variable plotted: TRMM daily precipitation rate (TRMM\_3B42\_daily\_precipitation\_V7)
- 6,574 daily data granules covering the globe at 0.25 deg resolution with latitude +/- 50 deg and date range: 1/1/1998 – 12/31/2015 (26 GB input data volume).
- Giovanni implementation uses highly optimized compiled code based on NetCDF Operator (NCO) toolkit, but is single threaded. NEXUS is implemented in Python and parallelized with Apache Spark.
- Giovanni execution time is compared with NEXUS for 16-way and 64-way parallelism.


# Time Averaged Map: Comparison of Giovanni and NEXUS Performance



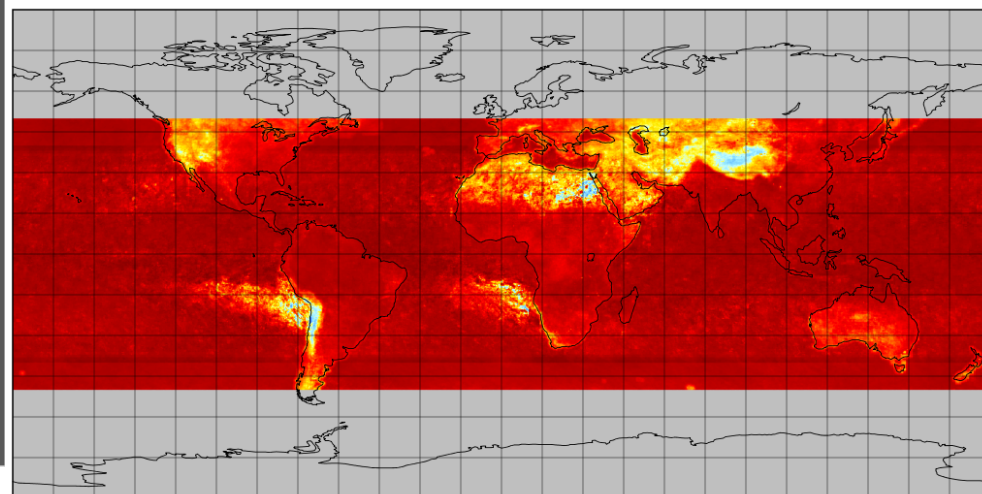
- NEXUS run on 8-node cluster computer at JPL running Solr, Cassandra, Spark 2.0, Mesos
- Global Time-Averaged Map
- Variable plotted: TRMM daily precipitation rate (TRMM\_3B42\_daily\_precipitation\_V7)
- 6,574 daily data granules covering the globe with latitude +/- 50 deg and date range: 1/1/1998 – 12/31/2015 (26 GB input data volume).
- Giovanni implementation uses highly optimized compiled code based on NetCDF Operator (NCO) toolkit, but is single threaded. NEXUS is implemented in Python and parallelized with Apache Spark.
- Giovanni execution time is compared with NEXUS for 16-way and 64-way parallelism.




# Correlation Map: Comparison of Giovanni and NEXUS Performance



Correlation of Precipitation Rate daily 0.25 deg. [TRMM TRMM 3B42 daily v7] mm/day vs. Ne...





Correlation of Precipitation Rate daily 0.25 deg. [TRMM TRMM 3B42 daily v7] mm/day vs. Near-Real-T...



- NEXUS run on 8-node cluster computer at JPL running Solr, Cassandra, Spark 2.0, Mesos
- Global Correlation Map
- Variable plotted: TRMM daily precipitation rate (TRMM\_3B42\_daily\_precipitation\_V7) vs TRMM real-time daily precipitation (TRMM\_3B42RT\_daily\_precipitation\_V7)
- 5,113 daily data granule pairs covering the globe with latitude +/- 50 deg and date range: 1/1/2001 – 12/31/2014 (40 GB input data volume).
- Giovanni implementation uses highly optimized compiled code based on NetCDF Operator (NCO) toolkit, but is single threaded. NEXUS is implemented in Python and parallelized with Apache Spark.
- Giovanni execution time is compared with NEXUS for 16-way and 64-way parallelism.

# Time Series: Comparison of Global vs. Subset Performance of Giovanni and NEXUS



- NEXUS run on 6 Amazon Web Services (AWS) Cloud instances of type “i2.4xlarge” running Solr, Cassandra, Spark 2.0, Mesos
- Area-Averaged Time Series over the indicated spatial subset (Global, State, City)
- Variable plotted: MODIS-Terra Aerosol Optical Depth (AOD) 550 nm dark target
- 5,789 daily data granules covering the globe at 1 deg resolution with date range: 3/1/2000 – 2/29/2016 (3 GB input data volume).
- Giovanni implementation uses highly optimized compiled code based on NetCDF Operator (NCO) toolkit, but is single threaded. NEXUS is implemented in Python and parallelized with Apache Spark.
- Giovanni execution time is compared with NEXUS for 16-way parallelism.
- NEXUS has superior performance for subsetting operations, as indicated by the ~300x speedup for the Colorado subset.



# Summary

## Benchmark NEXUS speedup factors over Giovanni

- ~300x speedup for 16-year area-averaged time series of Moderate Resolution Imaging Spectroradiometer (MODIS-Terra) Aerosol Optical Depth (AOD) at 1 degree resolution for Colorado with NEXUS running on 6 “i2.4xlarge” Amazon Web Services Cloud instances with Spark configured for 16-way parallelism.
- ~100x speedup for area-averaged time series of daily precipitation rate for the Tropical Rainfall Measuring Mission (TRMM with 0.25 degree spatial resolution) for the Continental United States over 18 years (1998 - 2015) with 64-way parallelism on an 8-node cluster computer at JPL.
- ~4x speedup for 18-year (1998 - 2015) TRMM daily precipitation global time averaged map (64-way parallel).
- ~22x speedup for 14-year (2001 - 2014) global map of correlation between TRMM daily and real time precipitation rate (64-way parallel).
- For small datasets, compiled, optimized, single- threaded executables like the NetCDF Operators (NCO) toolkit used in Giovanni work well.
- For large data analytics, NEXUS significantly outperforms due to its ability to horizontally scale in a cloud computing environment.

## Data tiling recommendations

- Tiling data into chunks yields significant performance benefits over monolithic global granule files, particularly for regional subsets.
- For calculations on small subsets, use a small tile size.
- For global or near-global calculations use larger tiles to optimize data read performance.

## JPL and NASA support open source development

- NEXUS is an open-source big data analytics framework, available at:  
<https://github.com/dataplumber/nexus>



# Jupyter Notebook and NEXUS

**Frank Greguska**

Jet Propulsion Laboratory  
California Institute of Technology

# Cloud Analysis Toolkit to Enable Earth Science (CATEES)

- **Goals**
  - Provide easy-to-use tools for working with EOSDIS data
  - Provide tools in a convenient package to users
  - Show users how to access EOSDIS API via script
  - Show users how to use analytics optimized cloud storage for analysis
  - Demonstrate cross-EOSDIS development
- **NEXUS contributions to CATEES**
  - Reproducibility: executable and sharable URLs
  - Client library integration with Jupyter Notebook
  - Using simple Python scripting to conduct analysis on large collection of EOSDIS data without having to download any file to local or remote server

## Index

**Module variables**

- ISO\_FORMAT
- session
- target

**Functions**

- daily\_difference\_average
- dataset\_list
- set\_target
- time\_series

**Classes**

- TimeSeries

## nexuscli.nexuscli module

Copyright (c) 2017 Jet Propulsion Laboratory, California Institute of Technology. All rights reserved.  
This module provides a native python client interface to the NEXUS (<https://github.com/dasapumber/nexus>) webservice API.

Usage:

```
import nexuscli
nexuscli.set_target('http://nexus-wohappi80031')
nexuscli.dataset_list()
```

**Module variables**

```
var ISO_FORMAT
var session
var target
```

**Functions**

```
def daily_difference_average(dataset, bounding_box, start_datetime, end_datetime)
Generate an anomaly Time series for a given dataset, bounding box, and timeframe.
dataset Name of the dataset as a string
bounding_box Bounding box for area of interest as a shapely.geometry.polygon.Polygon
start_datetime Start time as a datetime.datetime
end_datetime End time as a datetime.datetime
return List of TimeSeries namedtuples

def dataset_list()
Get a list of datasets and the start and end time for each.
return list of datasets. Each entry in the list contains shortname, start, end and
```

Jupyter 5 - Student Exercise - Answers Last Checkpoint: 16 minutes ago (unsaved changes)


```
In [3]: import time
import nexuscli
from datetime import datetime
from shapely.geometry import box

# TODO: Create a bounding box using the box method imported above
bbox = box(-150, 45, -120, 55)

# TODO: Plot the bounding box using the helper method plot_box
plot_box(bbox)

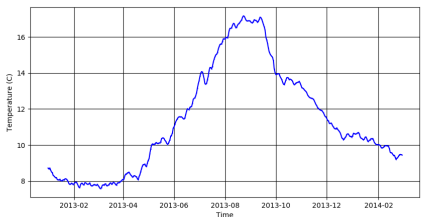
start = time.perf_counter()

# TODO: Call the time_series method for the AVHRR_OI_V4_GRSST_NCEP dataset using
# your bounding box and time period 2013-02-01 through 2014-02-01
datasets = ['AVHRR_OI_V4_GRSST_NCEP']
start_time = datetime(2013, 2, 1)
end_time = datetime(2014, 2, 1)
ts = nexuscli.time_series(datasets, bbox, start_time, end_time, sparse=True)
print('Time Series took {} seconds to generate'.format(time.perf_counter() - start))
```



Time Series took 66.89287395193062 seconds to generate

```
In [37]: # TODO: Plot the result using the 'show_plot' helper method
avhrr_ts = ts[0]
show_plot(avhrr_ts.time, avhrr_ts.mean, 'Time', 'Temperature (C)')
```





# BLANK



# Hands-On Labs

**Frank Greguska, Joseph Jacob, Nga Quach**

Jet Propulsion Laboratory  
California Institute of Technology