

National Aeronautics and
Space Administration

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

NEXUS

Big Data Analytics

Thomas Huang

NASA Physical Oceanography Distributed Active Archive Center (PO.DAAC)

Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Drive
Pasadena, CA 91109-8099
United States of America

© 2016 California Institute of Technology. Government sponsorship acknowledged.
Reference herein to any specific commercial product, process, or service by trade name,
trademark, manufacturer, or otherwise, does not constitute or imply its endorsements by the
United States Government or the Jet Propulsion Laboratory, California Institute of Technology.

THUANG/JPL

2016 ESIP Federation Summer Meeting



National Aeronautics and
Space Administration

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

“Traditional data management products are built on top of a scale-up architecture that cannot handle petabyte-scale applications...”

– *Why Traditional Data Management Fails in the Eras of NoSQL and Big Data*

“It’s not enough to just collect and store a lot of information; it’s what you do with the insights and learnings from that data that have real, measurable impacts on a business.”

– *The Art of Data Management*

“The Big Data technologies enable a whole new class of insight discovery, against data we’ve either not had access to before or not had the tools to efficiently process such as high volume time-series analysis or cluster analysis on a traditional RDBMS.”

– *Of Course Big Data Will Change the “Traditional” Data Warehouse*

“Climate change research is now ‘big science’, comparable in its magnitude, complexity and societal importance to human genomics and bioinformatics.”

– *Database bonanza*



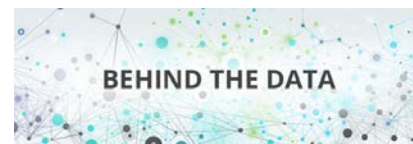
National Aeronautics and
Space Administration

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

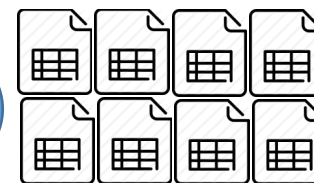
but ... How Do We Get There?



There

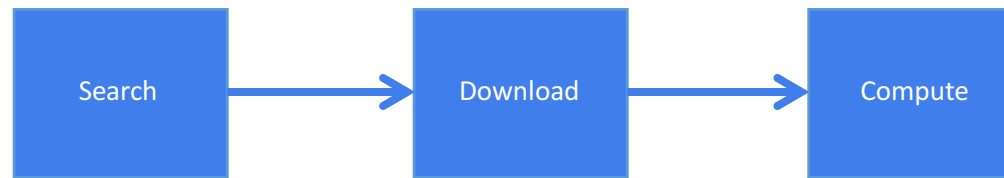


Here





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) yield poor performance
- 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



National Aeronautics and
Space Administration

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

NEXUS Deep Data Analytics: One-Minute Summary

NEXUS is an emerging technology developed at JPL

- A Cloud-based/Cluster-based data platform that performs scalable handling of observational parameters analysis designed to scale horizontally by
- Leveraging high-performance indexed, temporal, and geospatial search solution
- Breaks data products into small chunks and stores them in a Cloud-based data store

Data Volumes Exploding

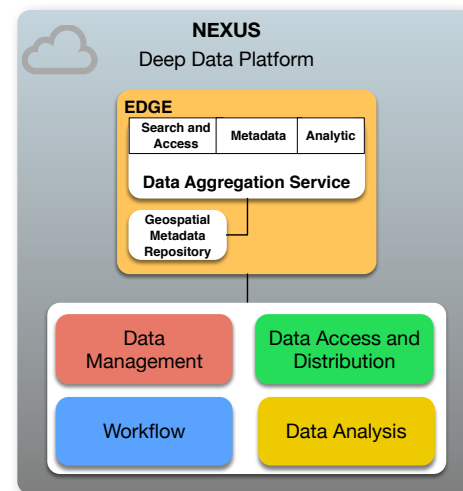
- SWOT mission is coming
- File I/O is slow

Scalable Store & Compute is Available

- NoSQL cluster databases
- Parallel compute, in-memory map-reduce
- Bring Compute to Highly-Accessible Data (using Hybrid Cloud)

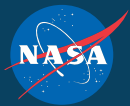
Pre-Chunk and Summarize Key Variables

- Easy statistics instantly (milliseconds)
- Harder statistics on-demand (in seconds)
- Visualize original data (layers) on a map quickly



Built with open source technologies

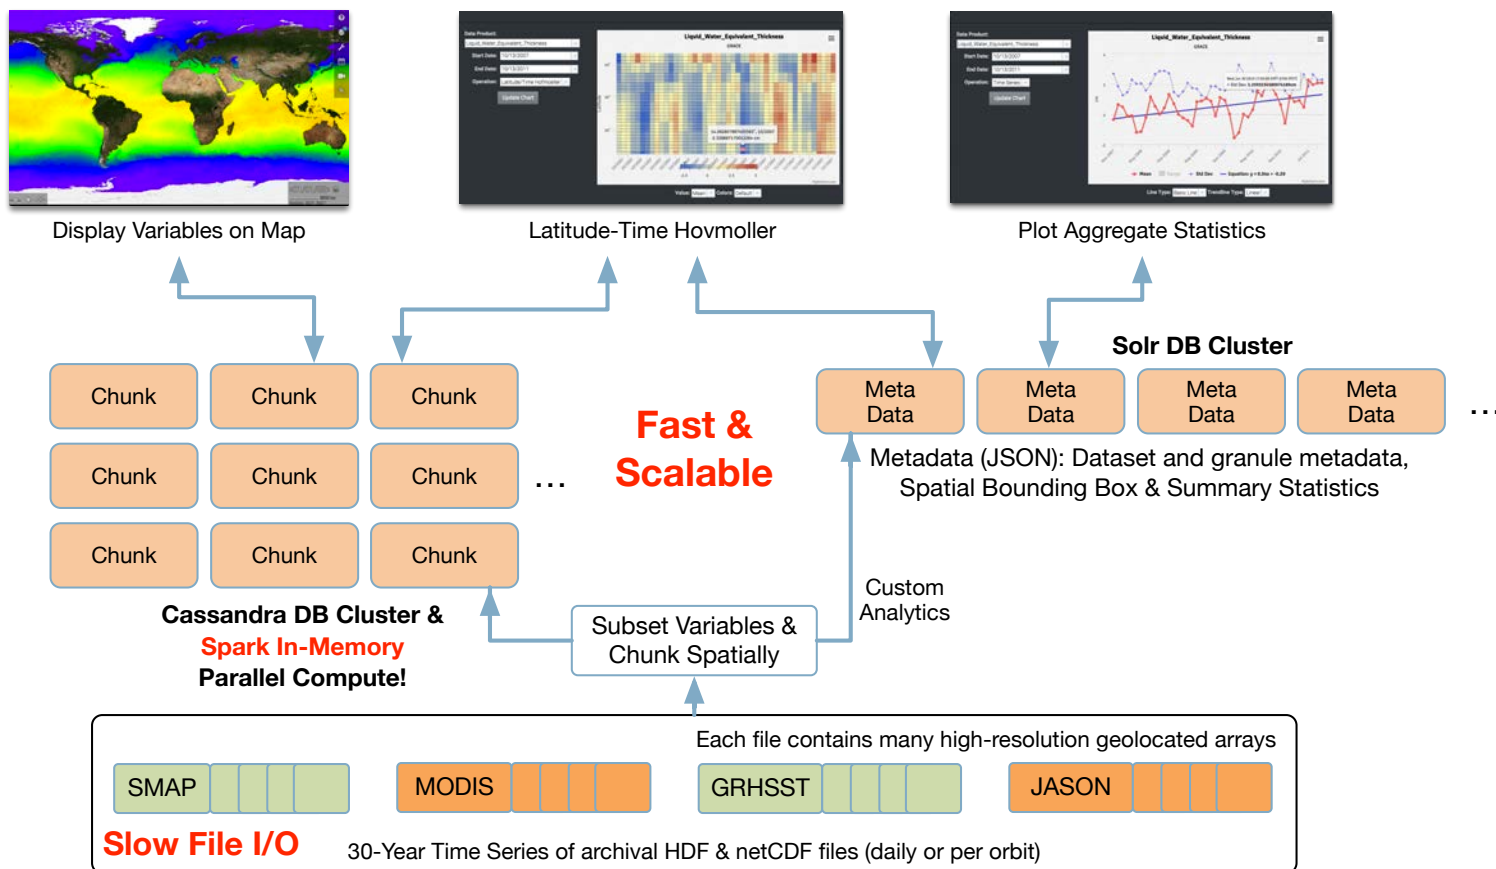
- Apache Solr
- Apache Cassandra
- Apache Spark/PySpark
- Apache Mesos/YARN
- Apache Kafka
- Apache Zookeeper
- Tornado
- Spring XD
- EDGE



National Aeronautics and
Space Administration

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

Analytics & Summarization Stack Using Two Scalable Database Technologies



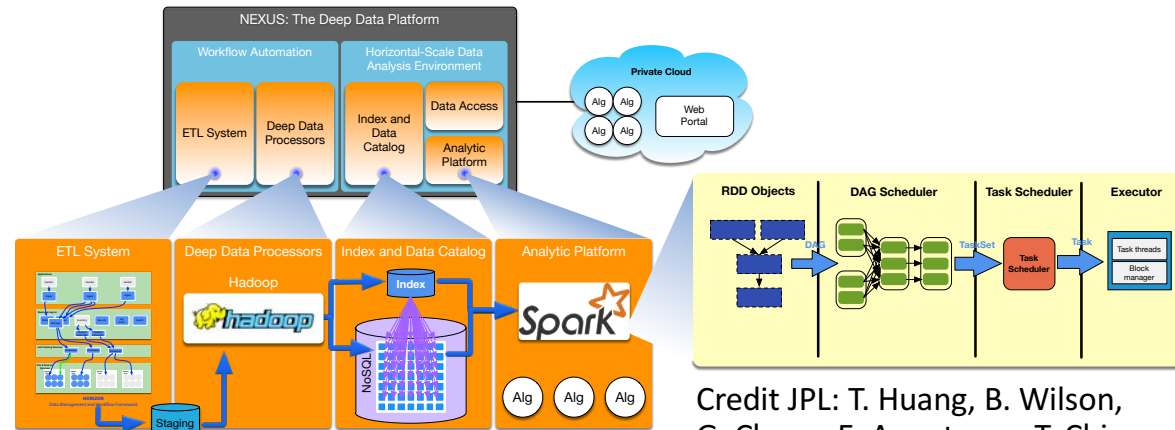


National Aeronautics and
Space Administration

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

Deep Data Computing Environment (DDCE)

- ETL System – Ingest and stage data
- Deep Data Processors – metadata, statistics, and tiles
- Index and Data Catalog – horizontal-scale geospatial search and tile retrieval
- Analytic Platform – Spark-based domain-specific analytics
- Data Access – tile and collection-based data access
- Cloud Platform – portal and custom VMs



Credit JPL: T. Huang, B. Wilson,
G. Chang, E. Armstrong, T. Chin

AIST-14: OceanXtremes



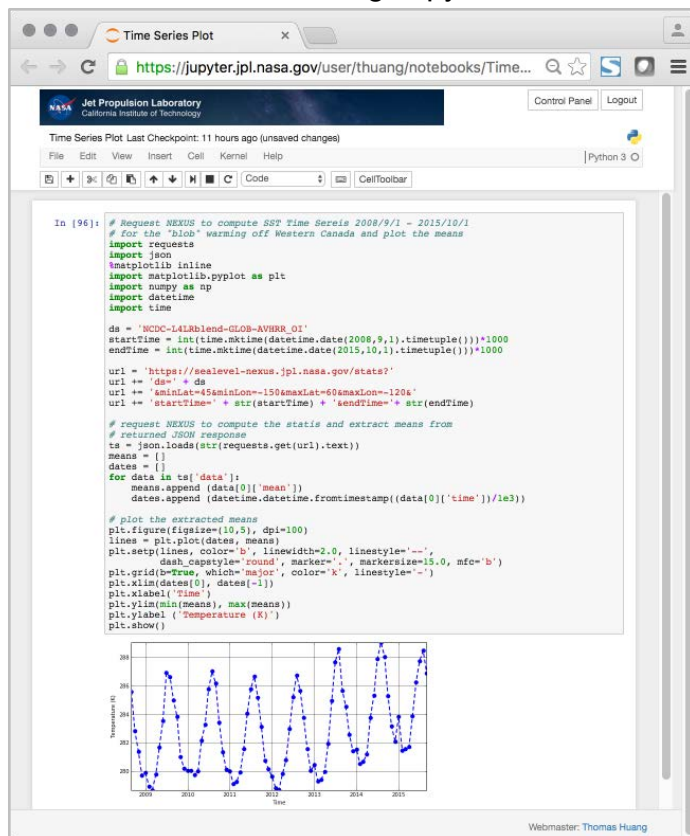


National Aeronautics and
Space Administration

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

The Notebook

Interact with NEXUS using Jupyter Notebook



- **/capabilities**: list of capabilities
- **/chunks**: list data chunks by location, time, and datasets
- **/correlationMap**: Correlation Map
- **/datainbounds**: Matchup operation to fetch values from dataset within geographic bounds
- **/datapoint**: Matchup operation to fetch value at lat/lon point
- **/dailydifferenceaverage**: Daily difference average
- **/latitudeTimeHofMoeller**: Latitude Time Hovmoeller
- **/list**: list available datasets
- **/longitudeLatitudeMap**: Longitude Latitude Map
- **/longitudeTimeHofMoeller**: Longitude Time Hovmoeller
- **/stats**: Statistics (standard deviation, count, min/max, time, mean)



National Aeronautics and
Space Administration

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

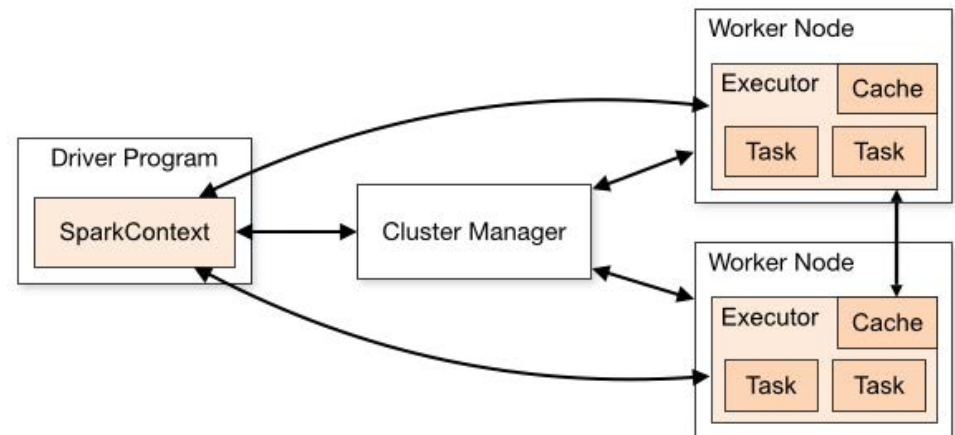
Apache Spark and Cassandra

- **Apache Spark**

- In-Memory Map-Reduce framework
- Datasets partitioned across a compute cluster by key
- Resilient Distributed Dataset (RDD)
 - Fault-tolerant, parallel data structures
 - Intermediate results persisted in memory
 - User controls the partitioning to optimize data placement
- Rich set of operators on RDD's: `Map`, `Filter`, `GroupByKey`, `ReduceByKey`, etc.
- Computation is implicit (Lazy) until answers needed
- Uses YARN/Mesos

- **Apache Cassandra**

- Horizontal-scale NoSQL database
- Constant-time writes regardless of the size of data set grows
- No-single-point of failure architecture





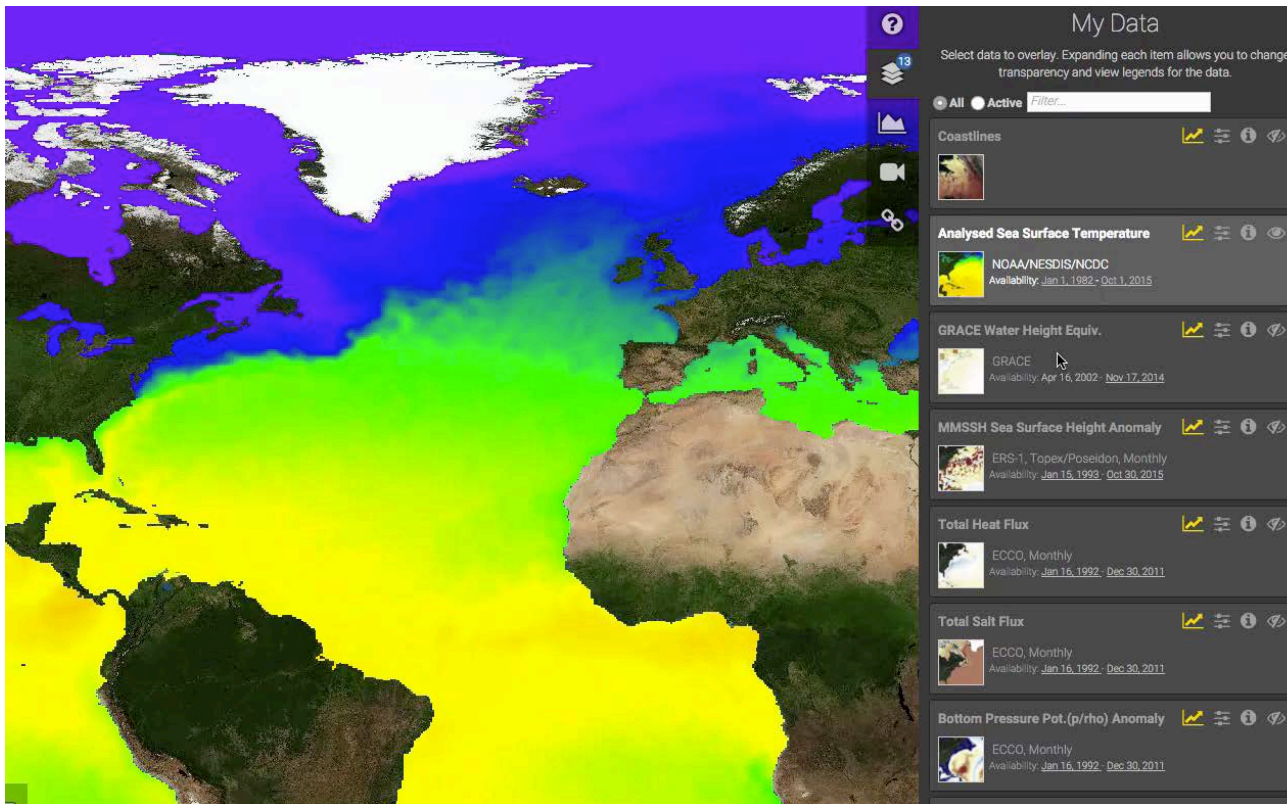
National Aeronautics and
Space Administration

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

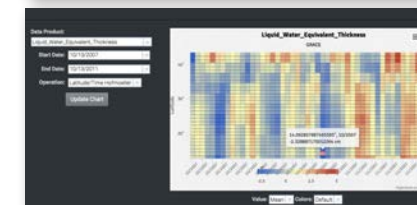
On-The-Fly Analysis for Sea Level Rise Research

<https://sealevel.nasa.gov>

Sea Level Change - Data Analysis Tool



- Visualizations
- Time Series
- Data Comparison
- Latitude/Time Hofmoeller
- Etc.



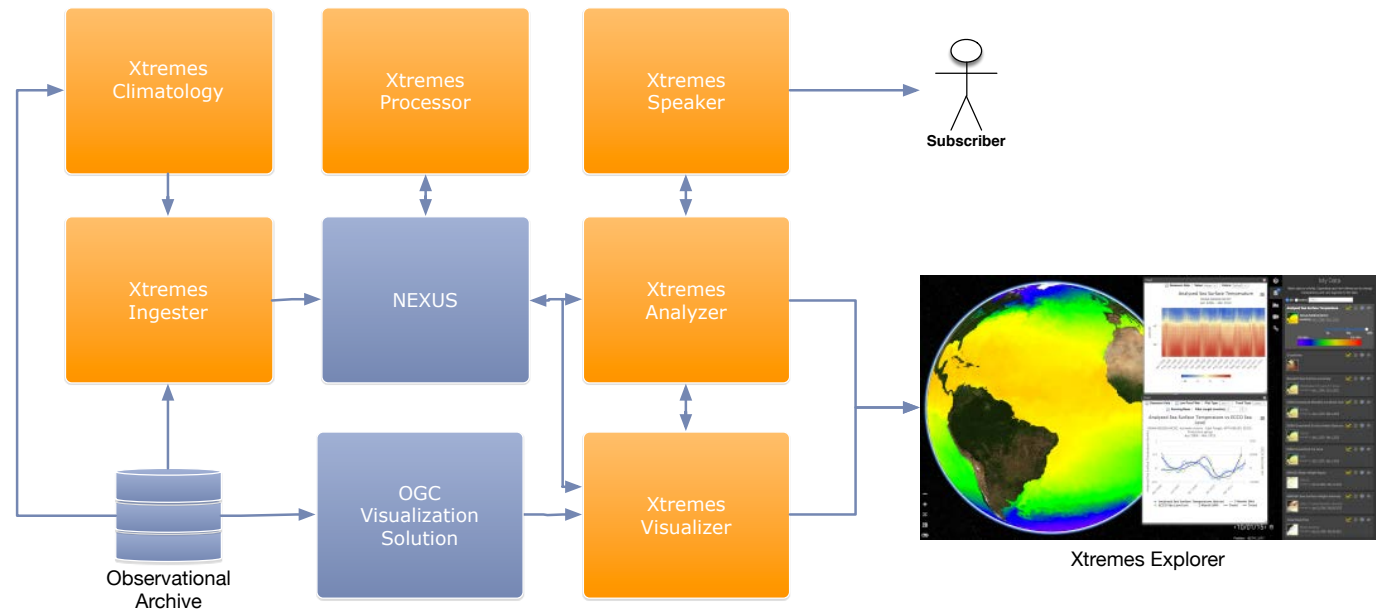


National Aeronautics and
Space Administration

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

AIST-14: OceanXtremes – Data-Intensive Anomaly Detection System

- Anomaly detection is a process of identifying items, events or observations outside the “norm” or expected patterns
- Current and future oceanographic missions and our research communities present us with challenges to rapidly identify features and anomalies in increasingly complex and voluminous observations
- Typically this is a two-stage procedure
 - Determine a long-term/periodic mean (“climatology”)
 - Deviations from the mean are searched. Step 1 could be omitted in cases where a climatology data set already exists.



OceanXtremes System Architecture

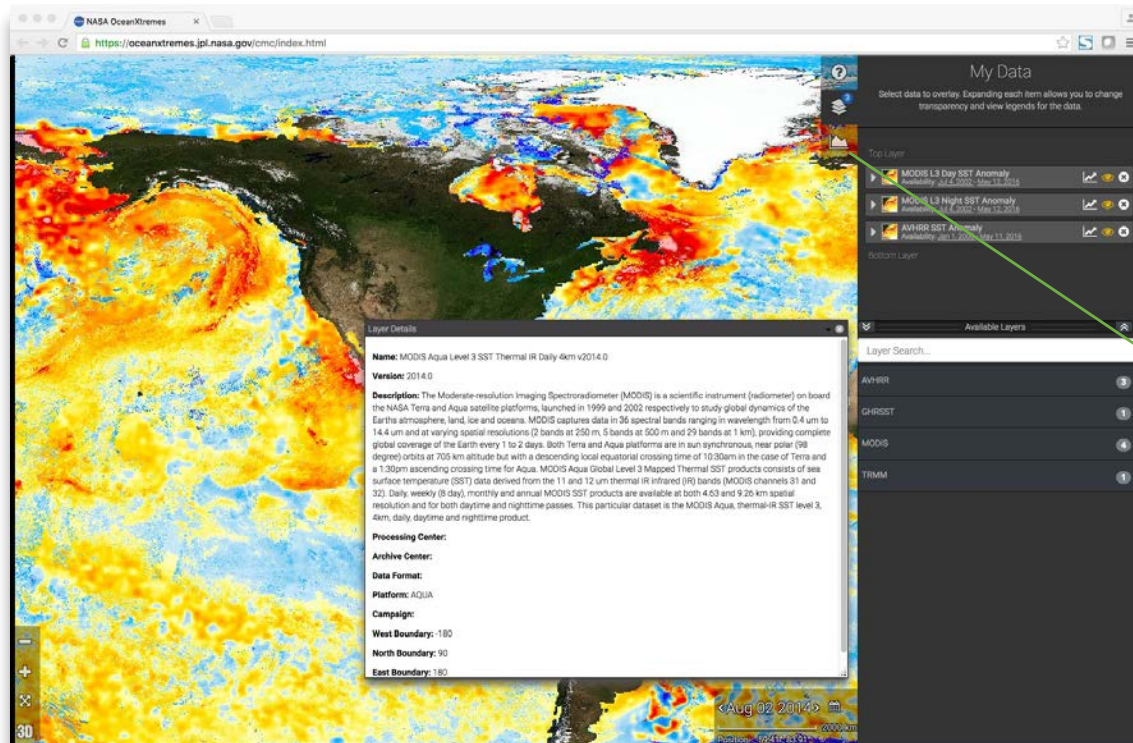
Credit: T. Huang, E. Armstrong, G. Chang, T. Chin, B. Wilson



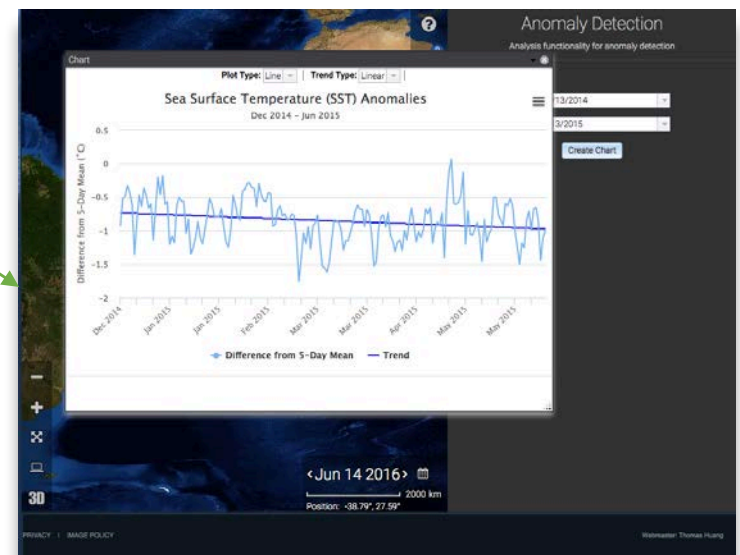
National Aeronautics and
Space Administration

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

Xtremes Explorer



High Resolution Data Visualization for the Web



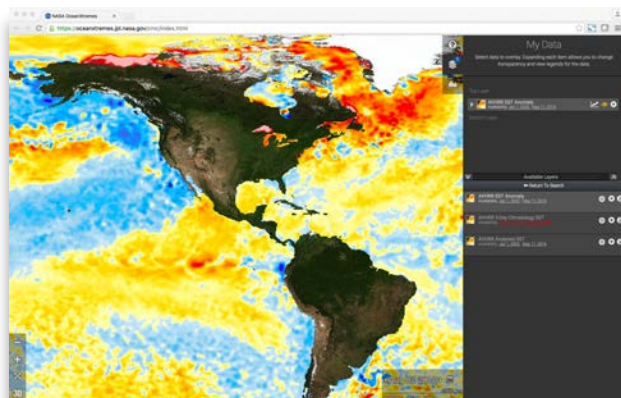
Data Analysis Workbench



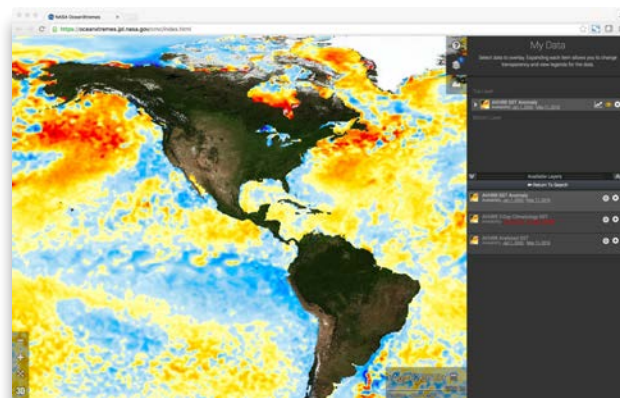
National Aeronautics and
Space Administration

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

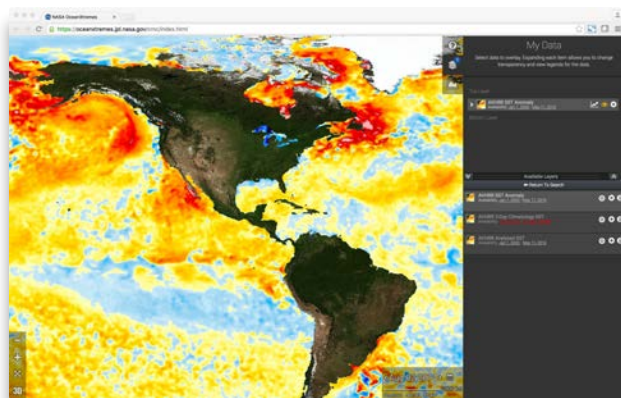
Daily Anomaly



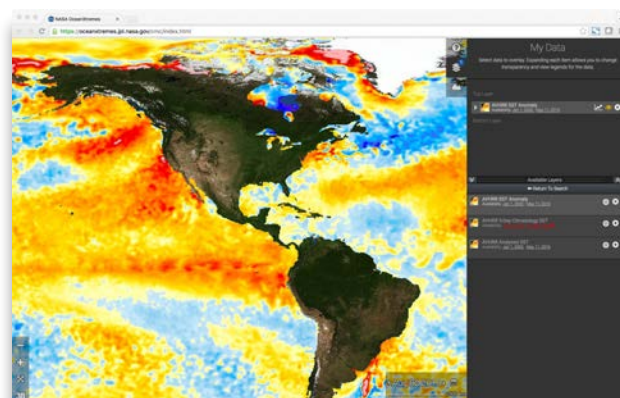
Aug 02, 2012



Aug 02, 2013



Aug 02, 2014



Aug 02, 2015

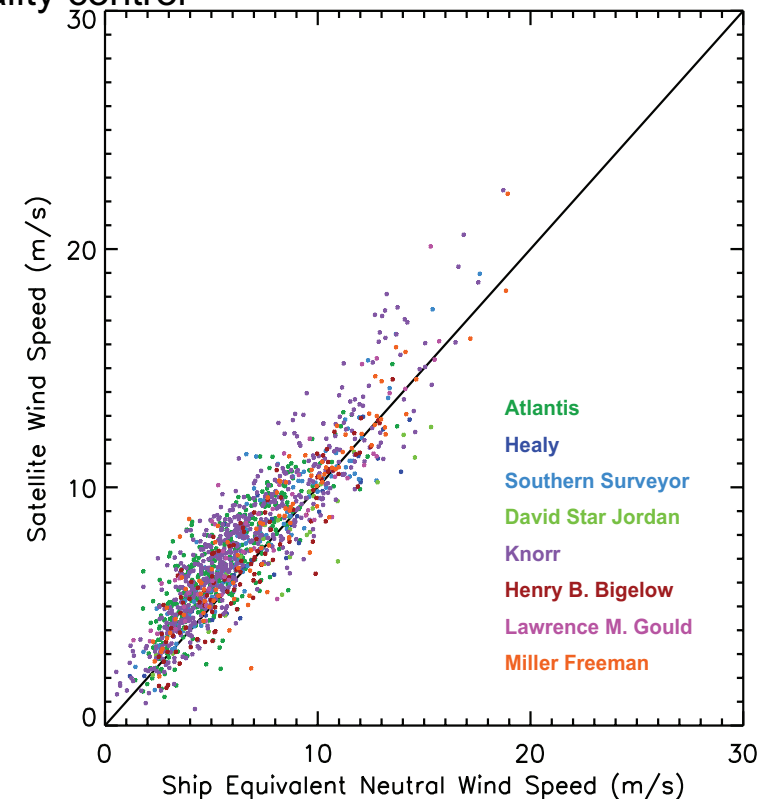


National Aeronautics and
Space Administration

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

AIST-14: Distributed Oceanographic Matchup Service (DOMS)

- A distributed data service (a.k.a DOMS) to match satellite and in situ marine observations to support platform comparisons, cross-calibration, validation, and quality control
- Use Cases
 - Satellite Cal/Val and algorithm development
 - Decision support
 - Planning field campaigns
 - Real-time operational activities
 - Scientific investigation
 - Process studies
 - Model assimilation services
 - User friendly interface to support student research
 - Alternate matching
 - Satellite to satellite
 - Satellite/in situ to model

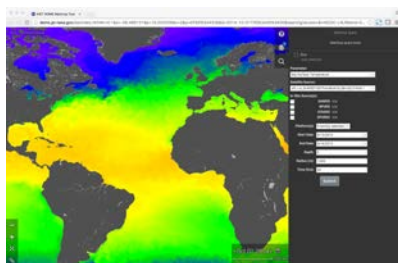




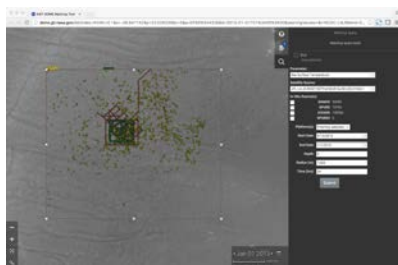
National Aeronautics and
Space Administration

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

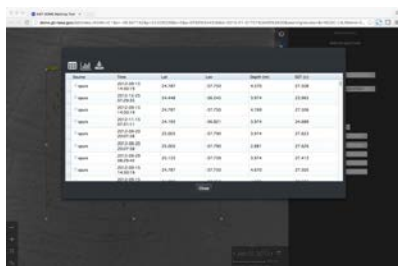
DOMS Architecture



High resolution
data visualization

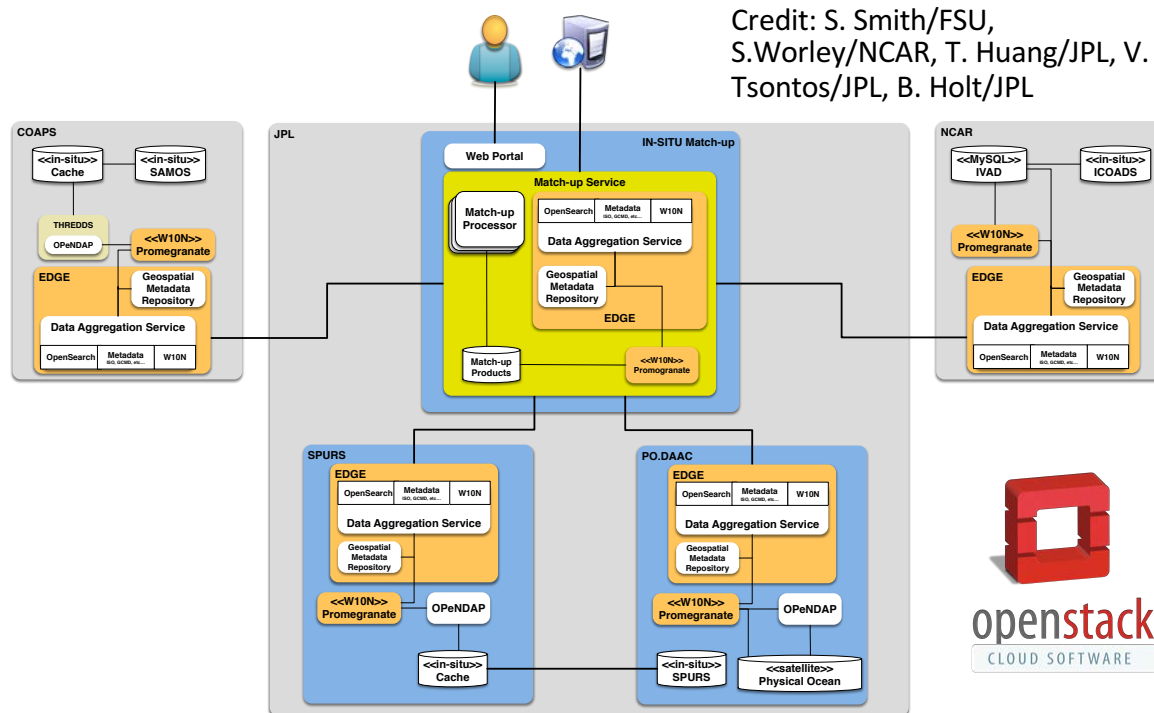


Initial match
according to user
initial selection



SST matchup with
SPURS

Dynamic Match-up



Credit: S. Smith/FSU,
S.Worley/NCAR, T. Huang/JPL, V.
Tsonetos/JPL, B. Holt/JPL



```
% curl
```

```
"https://doms.jpl.nasa.gov/doms/matchup2?s1=GHRSSST&s2=
spurs&s=2012-09-29T00:00:00Z&e=2012-09-
29T23:59:59Z&b=-38.74,23.93-
37.33,25.2&tt=3600&dt=50&rt=100"
```

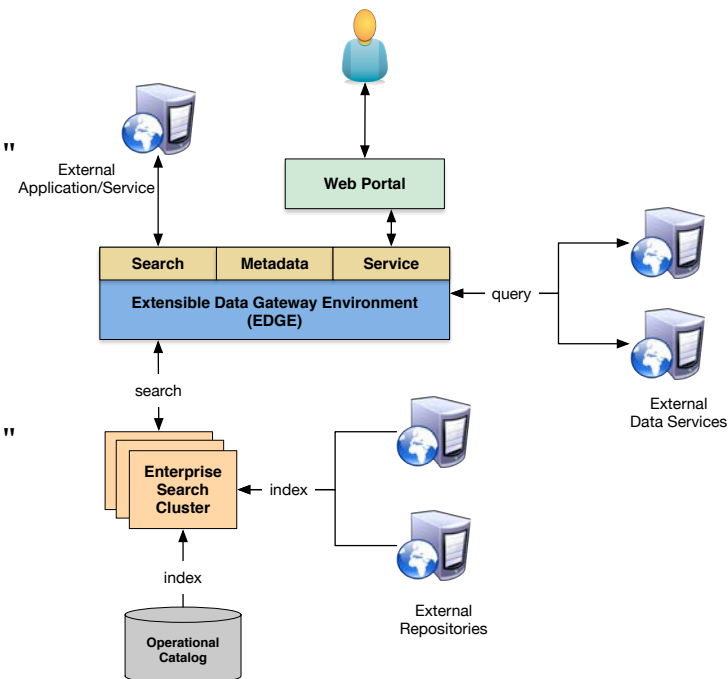


National Aeronautics and
Space Administration

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

Open Search for In-Situ Data

- **Common software interface for data search and retrieval (OpenSearch JSON)**
- **SAMOS @FSU**
 - Sample query: `curl "http://doms.coaps.fsu.edu/edge/samos?startTime=2012-08-01T00:00:00Z&endTime=2013-10-31T23:59:59Z&bbox=-45,15,-30,30"`
 - Integration with THREDDS and physical file archive
 - Full indexing of all relevant parameters: 54,911,530 documents (each with sst, wind, sss measurements)
- **ICOADS @NCAR**
 - Sample query: `curl "http://rda-db-icoads.ucar.edu:8890/ws/search/icoads?startTime=2012-08-01T00:00:00Z&endTime=2013-10-31T23:59:59Z&bbox=-45,15,-30,30"`
 - Integration with IVAD, MySQL server, hosts all ICOADS data exported from the original IMMA files
 - Full indexing of all relevant parameters: 517,140,830 documents (each with sst, wind, sss measurements)
- **SPURS Data Service**
 - Sample query: `curl "https://doms.jpl.nasa.gov/ws/search/spurs?startTime=2012-08-01T00:00:00Z&endTime=2013-10-31T23:59:59Z&bbox=-45,15,-30,30"`
 - 13,817,392 documents (each with sst, wind, and sss measurements)



Open Source, available on Github.com
<https://github.com/dataplumber/edge>

NEXUS Performance Challenge

The Challenge: Show that NEXUS performs 2X or greater speed improvement compare to Giovanni – Sponsored by NASA/ESDIS

- Dataset: TRMM Daily Precipitation (TRMM 3B42 Daily V7), 18 years, 6574 granule files, ~26GB
- Giovanni – backed by the popular NCO library
 - Highly optimized C/C++ library
- NEXUS
 - Apache Solr for spatial searches, metadata, and pre-computed statistics
 - Apache Cassandra for clustered data storage where granule data is partitioned into tiles
 - Apache Spark for data analytic platform

Question: With smaller tiles enables more parallelism. Does more parallelism yield faster performance?

- Divide daily granule into 16 tiles – 16 Executors and 16 Data Partitions
- Also experimented with smaller tiles - 545 tiles – 64 Executors and 128 Data Partitions

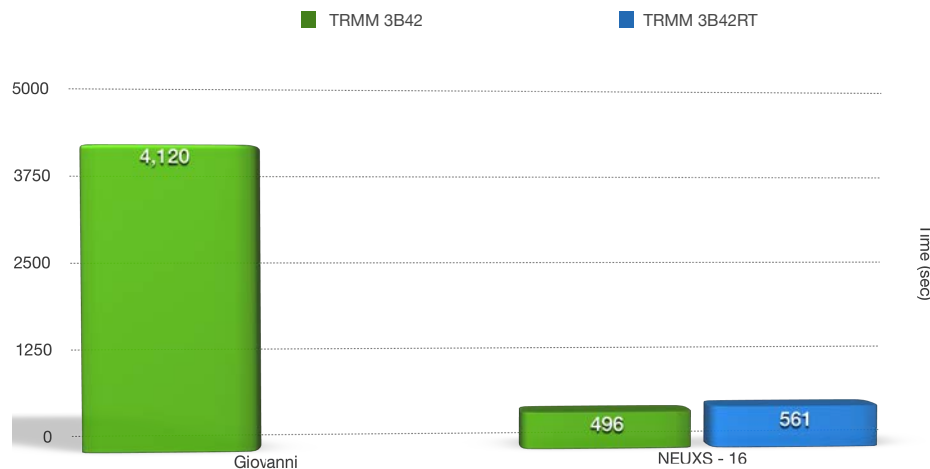


National Aeronautics and
Space Administration

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

NEXUS Performance Challenge Area Averaged Time Series

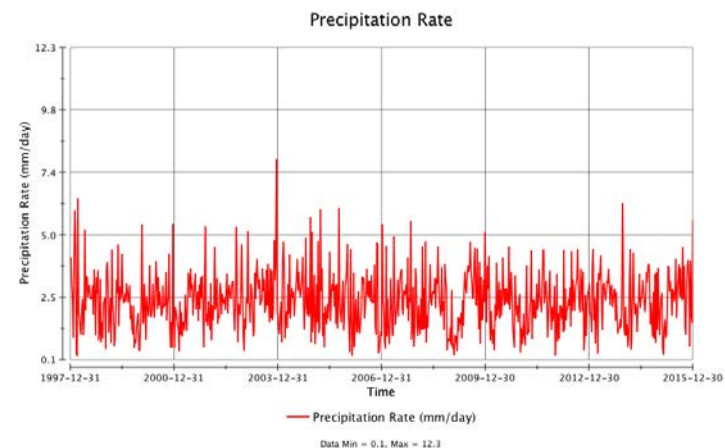
Area Averaged Time Series: The standard time-series plot produced by computing spatial averages over the user-selected area of a given variable for each time step within the user's range. Fill values do not contribute to the spatial averages. Each average value is then plotted against time to create the time-series output.



Global 18-year Area Averaged Map - Performance

Observation

- On-the-fly data subsetting
- Giovanni took over an hour
- NEXUS, with smaller tiles, took less than 5min
- More tiles and more executors yield FASTER performance
- With 16-way parallel – NEXUS performs ~8X faster
- **With 64-way parallel – NEXUS performs ~15X faster**



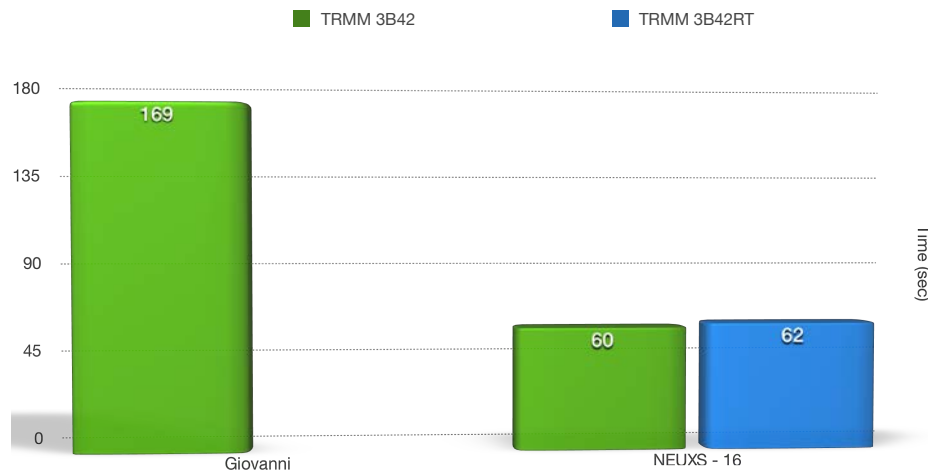


National Aeronautics and
Space Administration

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

NEXUS Performance Challenge Global 18-Year Time Averaged Map

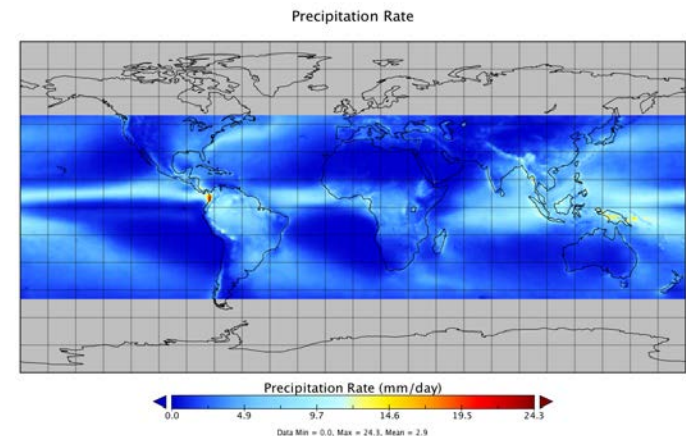
Time Averaged Map: The time averaged map shows data values for each grid cell within the user-specified area, averaged (linearly) over the user-specified time range as a map layer. Fill values do not contribute to the time averages.



Global 18-year Time Averaged Map - Performance

Observation

- Global coverage
- Giovanni took about 3min
- NEXUS, with bigger tiles, took about 1min
- More tiles and more executors yield SLOWER performance
- With 16-way parallel – NEXUS performs ~2.8X faster



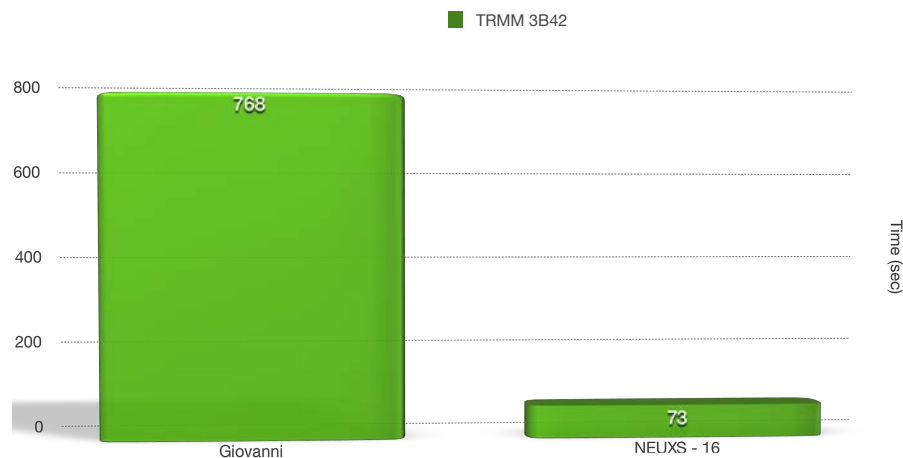


National Aeronautics and
Space Administration

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

NEXUS Performance Challenge Global 14-Year Correlation Map

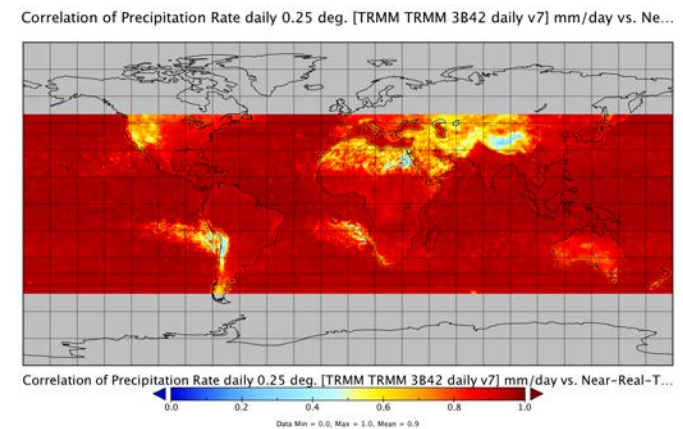
Global Correlation Map: The correlation map calculates correlation coefficient using simple linear regression between two variables over time within each grid cell, producing two maps: one showing the correlation coefficient (R) and the other displaying the number of contributing (matching) samples in each grid cell. Any grid cell that contains fewer than three matched pairs over time will be assigned a fill value.



Global 14-year Correlation Map - Performance

Observation

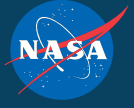
- Global coverage
- Giovanni took over 10min
- NEXUS, with bigger tiles, took a little over 1min
- More tiles and more executors yield SLOWER performance
- With 16-way parallel – NEXUS performs ~7X faster





Key Takeaways

- Global analysis - NEXUS out performs Giovanni with larger tiles and limit parallelism
- Area analysis - NEXUS out performs Giovanni with finer tiling scheme
- Another observation: NEXUS is a fast data subsetter
- More parallelism \neq Faster performance
 - For global analytics, more smaller tiles actually yields slower performance
 - Why? Scheduling, data transport, data queries, etc.
- Big Data \neq Cloud Computing
- Big Data is not a new computing problem. Cloud Computing opens up new approaches in tackling Big Data
- While Cloud Computing has many benefits, it only plays a part in the overall Big Data architecture
 - Process, Information Model, Technologies, etc.
- Apply Cloud Computing where it make senses
 - Data-Intensive Science, Cost reduction, Service reliability, etc.
- Part of the architectural design involves modernizing existing software solutions in order to
 - Truly leverage the elasticity of the Cloud
- It makes sense to bring the computing close to the data - On-premise Cloud (currently)
 - Need local experts
 - Governance
- Use automation deployment – Puppet, Chef, Salt
- Many technologies are mature in their standalone context
 - It doesn't mean they are high TRL when integrated into our domain-specific architecture
- Look into Open Source Solutions before build your own
- **NEXUS Current Activities**
 - It is still actively being developed and more to come
 - Part of the AIST Data Container Study



National Aeronautics and
Space Administration

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

July 14, 2016

**** APPROVED ****

The software known as NEXUS: Deep Data Platform (NTR-50157) has been approved for release as open source. You are authorized to upload the software to open source repository when you are ready to do so.



National Aeronautics and
Space Administration

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

IN008: Big Data Analytics



<https://agu.confex.com/agu/fm16/preliminaryview.cgi/Session13144>

IN008: Big Data Analytics

[Submit an Abstract to this Session](#)

Session ID#: 13144

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, and visualization. This session is to capture the latest on applying Cloud Computing for Big Data Analytical problems in all Earth and space domains. Topics include demonstration, studies, methods, solutions and solution discussion on

- Solutions for big data analytics
- Big data 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 **Christopher Lynnes**, NASA Goddard Space Flight Center, Greenbelt, MD, United States

Index Terms:

1914 Data mining [INFORMATICS]
1918 Decision analysis [INFORMATICS]
1926 Geospatial [INFORMATICS]
1976 Software tools and services [INFORMATICS]

Abstracts Submitted to this Session:

[Sharing Planetary-Scale Data in the Cloud \(Invited\)](#) (120851)

Joe Flasher, Amazon Web Services, Inc, Washington, DC, United States and **Jed Sundwall**, Amazon Web Services, Global Open Data Lead, Seattle, WA, United States

See more of: [Earth and Space Science Informatics](#)



National Aeronautics and
Space Administration

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

Special Thanks

NASA

Mike Little
Chris Lynnes
Kevin Murphy

FSU COAPS

Mark Bourassa
Jocelyn Elya
Shawn Smith
Adam Stallard

NCAR

Steve Worley
Ji Zaihua

JPL NEXUS Engineers

Kevin Gill
Frank Greguska
Joseph Jacob
Nga Quach
Brian Wilson

JPL Science Contributors

Ed Armstrong
Andrew Bingham
Carmen Boening
Mike Chin
Ben Holt
Rob Toaz
Vardis Tsontos
Victor Zlotnicki

Questions, and more information

Thomas.Huang@jpl.nasa.gov