National Aeronautics and
Space Administration

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

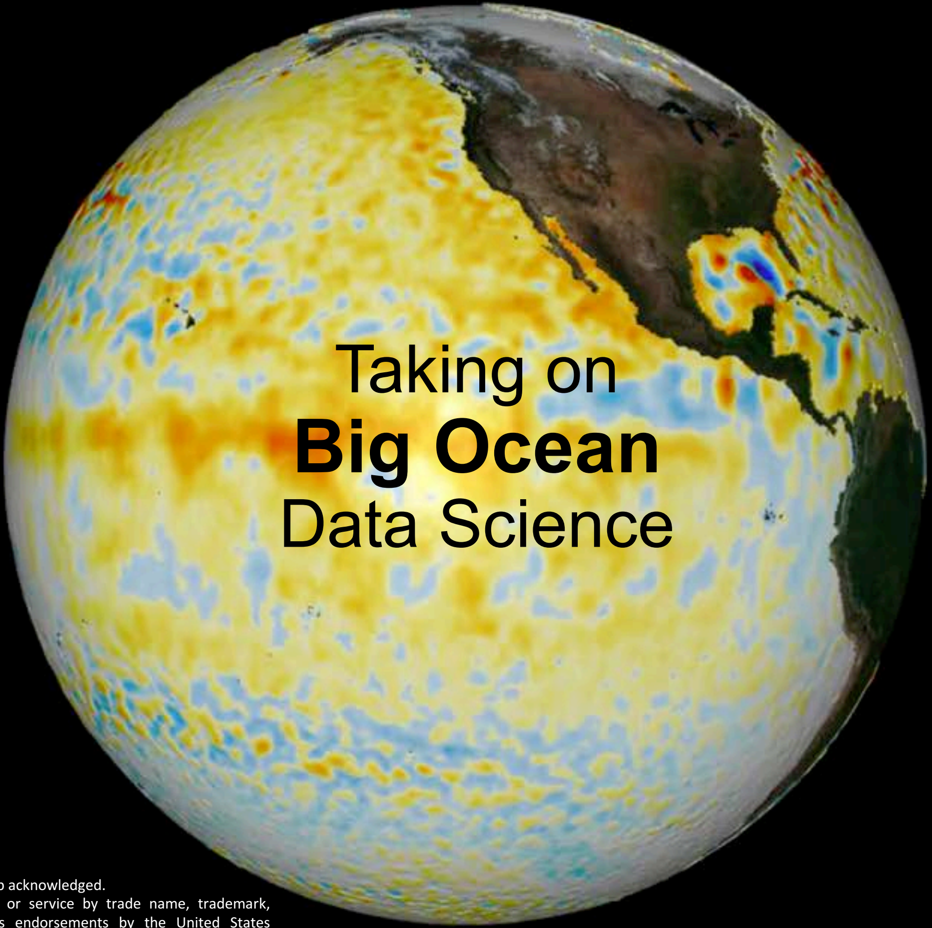

Taking on **Big Ocean** Data Science

Thomas Huang

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

CL#16-4603

IMDIS 2016, Gdansk, Poland

Data Scientist @NASA/JPL

- **Project Technologist** for the NASA's Physical Oceanography Distributed Active Archive Center (PO.DAAC) – http://podaac.jpl.nasa.gov
- **Architect** for the NASA Sea Level Change Portal – https://sealevel.nasa.gov
- **Principal Investigator / Co-Investigator** in several NASA-funded Big Data Analytic Projects
  - OceanXtremes: Oceanographic Data-Intensive Anomaly Detection and Analysis Portal – https://oceanxtremes.jpl.nasa.gov
  - Distributed Oceanographic Matchup Service (DOMS) – https://doms.jpl.nasa.gov
  - Mining and Utilizing Dataset Relevancy from Oceanographic Datasets (MUDROD)
  - Enhanced Quality Screening for Earth Science Data – https://vqss.jpl.nasa.gov
  - NEXUS - Big Data Analytic on the Cloud
- **Architect** for Tactical Data Science Framework for Naval Research
- **Ontologist** for the Semantic Web for Earth and Environmental Terminology (SWEET) Ontologies – http://sweet.jpl.nasa.gov
- **Chair** for The Federation of Earth Science Information Partners (ESIP) Cloud Computing Cluster
- **Chair/Co-Chair** for the NASA Earth Science Data System Working Groups

NASA'S PO.DAAC

- The **NASA Physical Oceanography Distributed Active Archive Center (PO.DAAC)** at Jet Propulsion Laboratory is an element of the **Earth Observing System Data and Information System (EOSDIS)**. The EOSDIS provides science data to a wide communities of user for NASA's Science Mission Directorate.

- Archives and distributes data relevant to the physical state of the ocean

- The mission of the PO.DAAC is to **PRESERVE** NASA's ocean and climate data and make these universally **ACCESSIBLE** and **MEANINGFUL**.

http://podaac.jpl.nasa.gov

Big Data Reality

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 (simulations, distributed linear algebra). Efficiency matters.

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

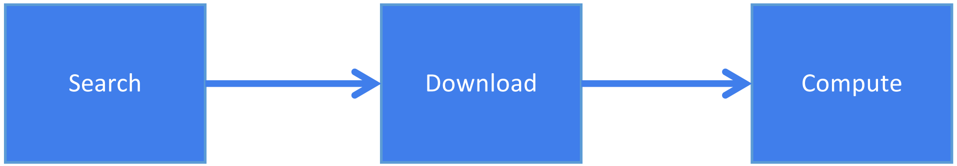

- Depending on the data volume (size and number of files)
- It could take many hours of download – (e.g. 10yr of observational data could yield thousands of files)
- It could take many hours of computation
- It requires expensive local computing resource (CPU + RAM + Storage)
- After result is produced,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

National Aeronautics and Space Administration

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

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**

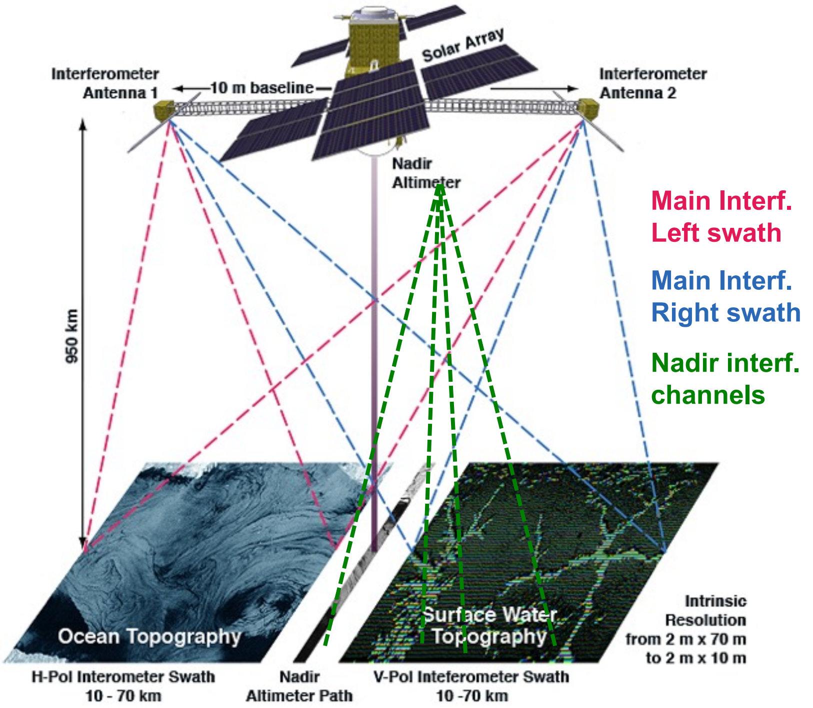

Launches April of 2021
https://swot.jpl.nasa.gov

The Silver Bullet?

Moore's Law is coming to an end due to physical limits of CMOS
"the number of transistors can put on a microchip doubles every year or so."

Cloud Computing provides an elastic infrastructural approach to Big Data

But… How Do We Get There?

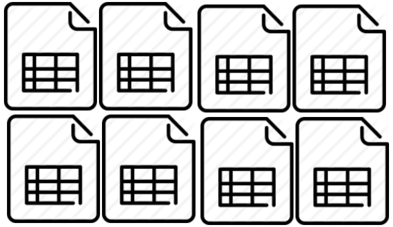

National Aeronautics and Space Administration
Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

NEXUS Deep Data Analytics: One-Minute Summary

NEXUS is an emerging technology developed at JPL

- Open Source: https://github.com/dataplumber/nexus
- 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

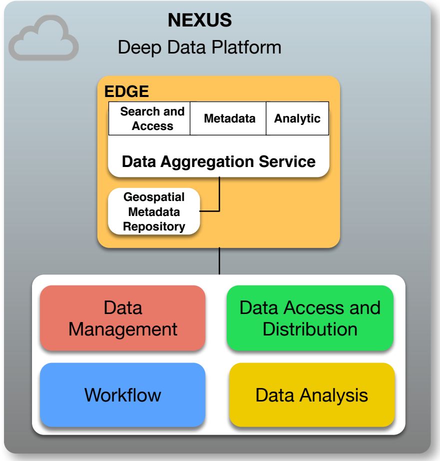

Built with open source technologies

- 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

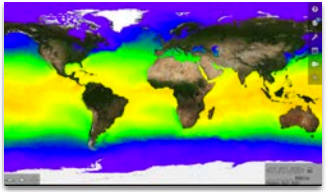

Display Variables on Map

Latitude-Time Hovmoller

Plot Aggregate Statistics

Solr DB Cluster

Meta Data | Meta Data | Meta Data | Meta Data …

Metadata (JSON): Dataset and granule metadata, Spatial Bounding Box & Summary Statistics

Chunk | Chunk | Chunk
Chunk | Chunk | Chunk …
Chunk | Chunk | Chunk

Fast & Scalable

Cassandra DB Cluster &
Spark In-Memory
Parallel Compute!

Subset Variables & Chunk Spatially

Custom Analytics

Each file contains many high-resolution geolocated arrays

SMAP | MODIS | GRHSST | JASON

Slow File I/O

30-Year Time Series of archival HDF & netCDF files (daily or per orbit)

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

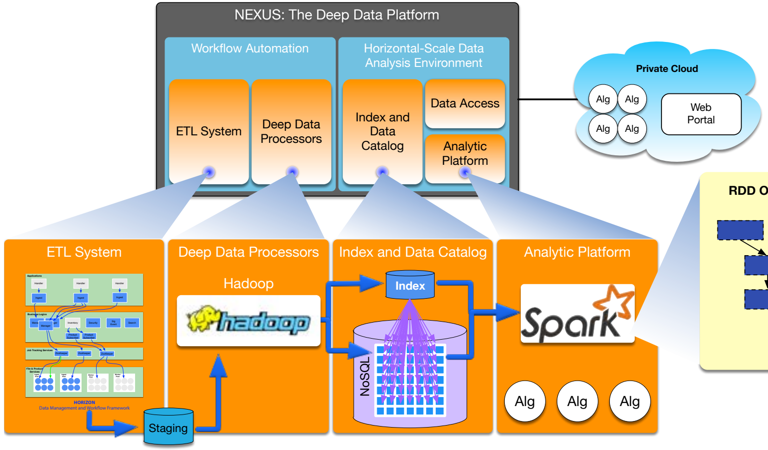

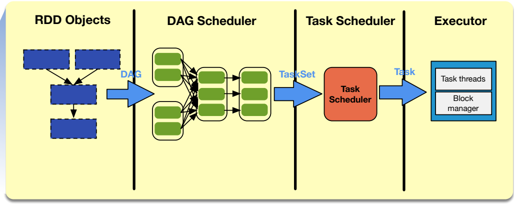

Credit JPL: T. Huang, B. Wilson, G. Chang, E. Armstrong, T. Chin

AIST-14: OceanXtremes

Using 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

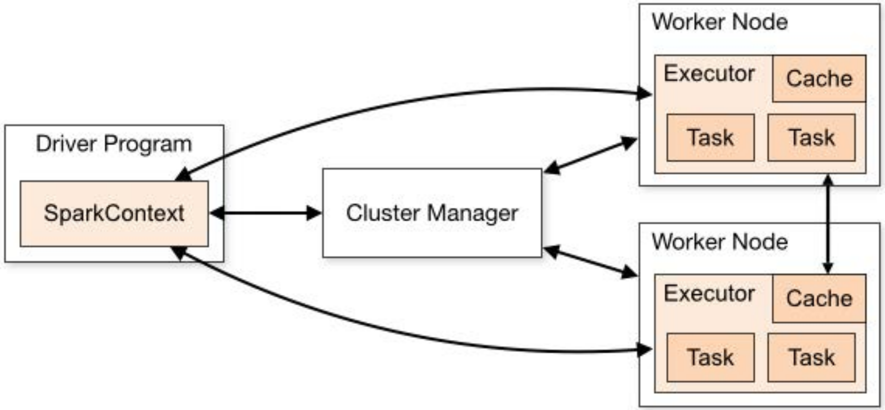

National Aeronautics and Space Administration

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

NEXUS Real Time Analysis

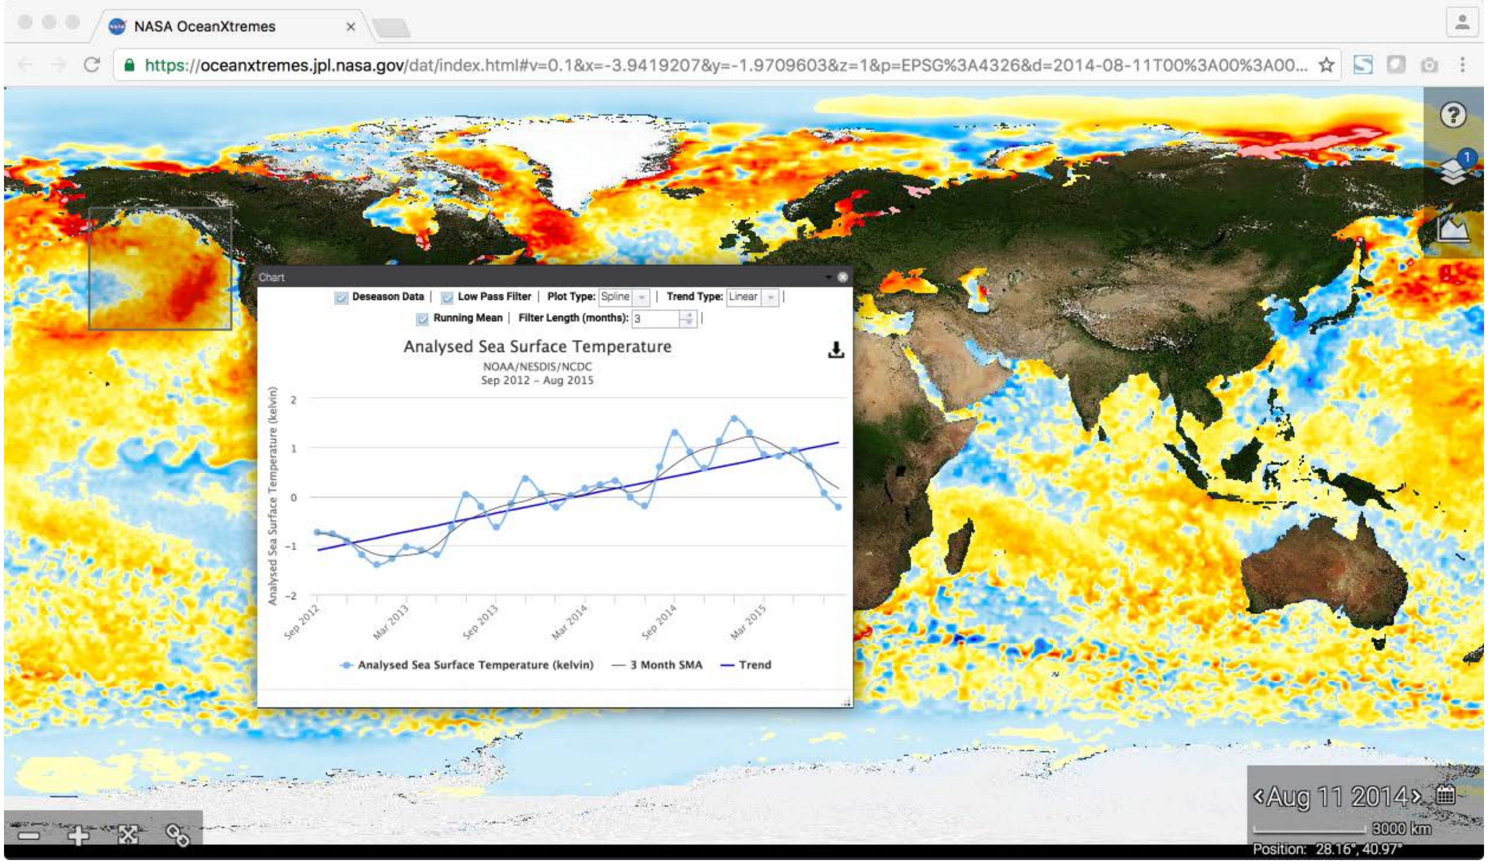

The Notebook

National Aeronautics and
Space Administration

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

Interact with NEXUS using Jupyter Notebook

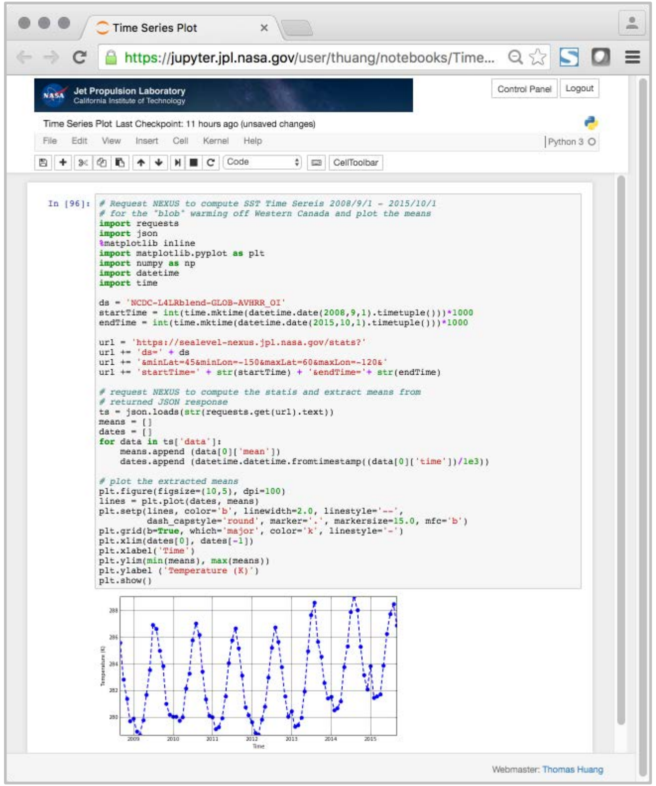

- **`/capabilities`**: list of capabilities
- **`/chunks`**: list data chunks by location, time, and datasets
- **`/correlationMap`**: Correlation Map
- **`/datainbounds`**: Matchup operation to fetch values from dataset within geographic bounds
- **`/datapoint`**: Matchup operation to fetch value at 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)

NEXUS 2.0 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
- **Algorithms**
  - Area Averaged Time Series
  - Global Time Averaged Map
  - Correlation Map
- **Giovanni** – web application for researchers to analyze NASA's gridded data.
  - Backed by the popular NCO (NetCDF Operator) 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

Performance Statistics

All performed on Apache Spark cluster with 16-way parallel

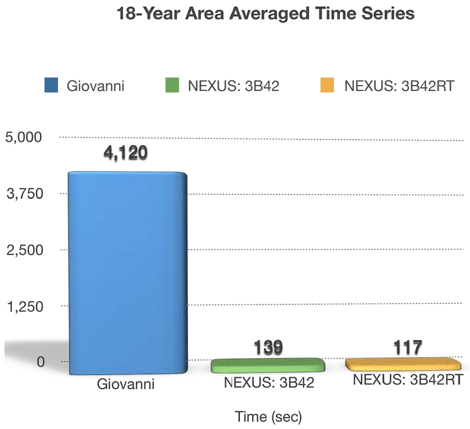

Giovanni: over an hour
NEXUS: a little over 2min
~30X faster

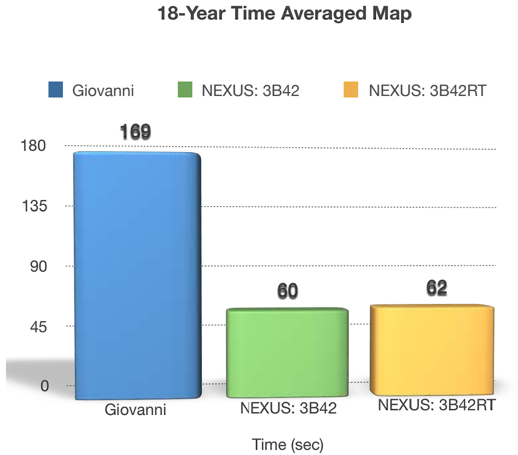

Giovanni: about 3min
NEXUS: 1min
~3X faster

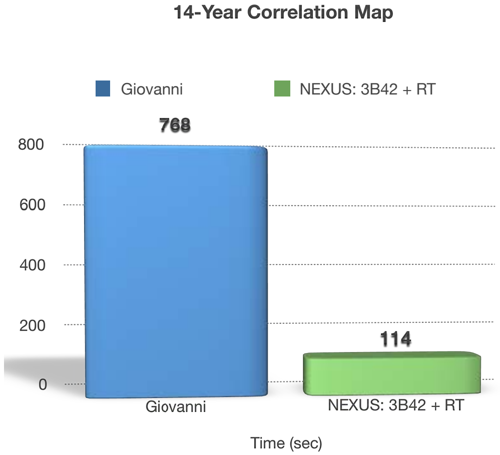

Giovanni: about 13min
NEXUS: 2min
~7X faster

National Aeronautics and Space Administration

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

NASA Sea Level Change Portal

Goals

- Provide scientists and the general public with a “one-stop” source for current sea level change information and data
- Provide interactive tools for accessing and viewing regional data
- Provide virtual dashboard for sea level indicators
- Provide latest news, quarterly report, and publications
- Provide ongoing updates through a suite of editorial products
  - Content articles
  - Multimedia

Features

- Featured news
- Sea level indicators
- Understanding sea level
  - Causes
  - Observations
  - Projections
  - Adaptation
- Data search
- Data Analysis Tool
- Ice Sheet Simulation and Modeling Tool
- Multimedia
- Sea level news
- Scientist interviews
- Publications
- Commentary
- Featured multimedia
- Subscription for newsletter

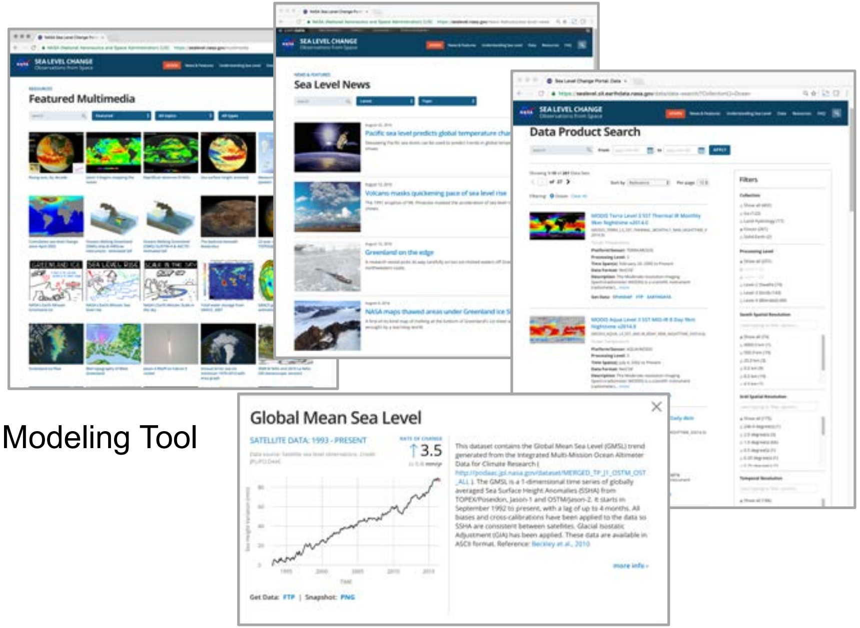

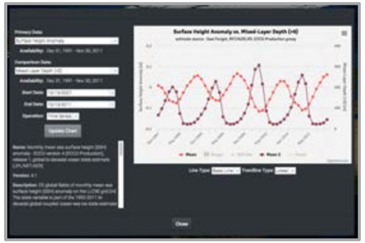

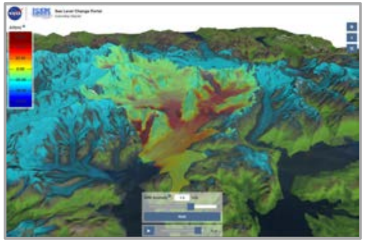

https://sealevel.nasa.gov

National Aeronautics and Space Administration

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

On-The-Fly Analysis for Sea Level Rise Research

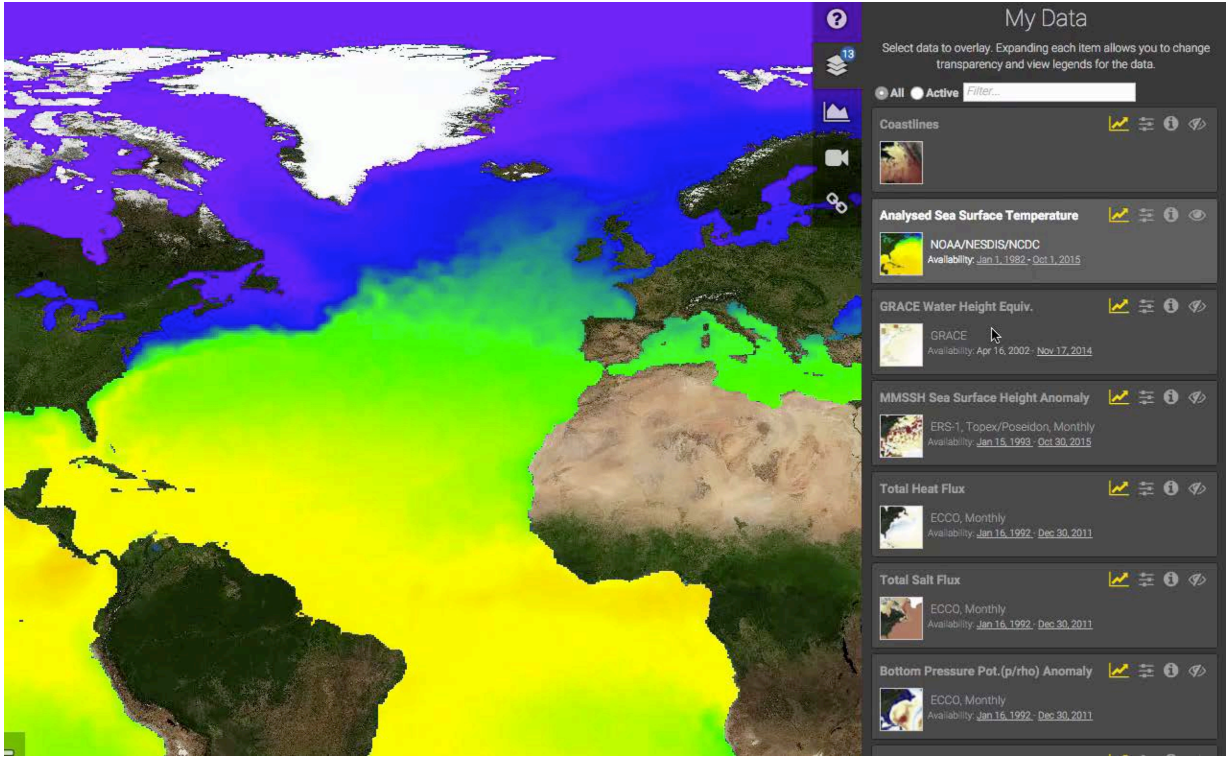

Sea Level Change - Data Analysis Tool

- Visualizations – WMTS – tiled imagery webserivce
- Time Series
- Data Comparison
- Latitude/Time Hofmoeller
- Etc.

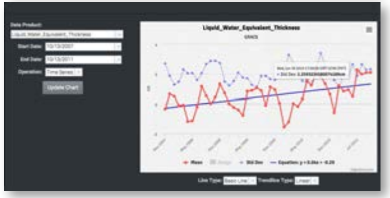

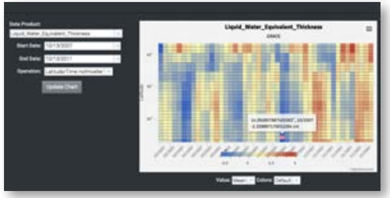

Ocean Anomalies

Identifying observations which do not conform to an expected pattern in a dataset or time series.

OceanXtremes – Data-Intensive Anomaly Detection System

- Funded by the NASA Advanced Information System Technology program
- 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
- Goals
  - Provide one-stop portal registry of ocean anomalies
  - Provide on-the-fly analysis and mining on observational data
- Typically this is a two-stage procedure
  1. Determine a long-term/periodic mean ("climatology")
  2. Deviations from the mean are searched. Step 1 could be omitted in cases where a climatology data set already exists.

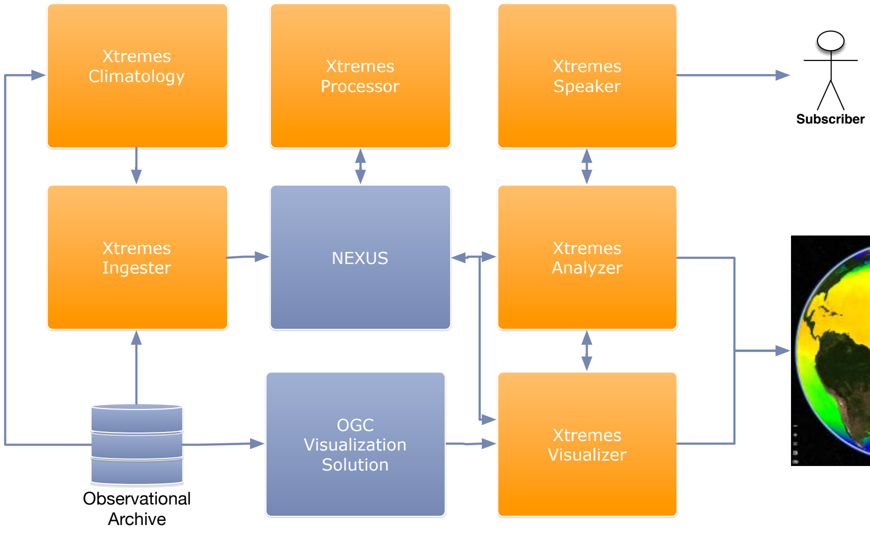

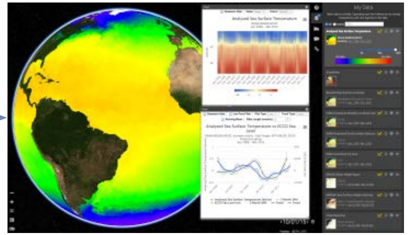

Xtremes Explorer

OceamXtremes 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: Daily Anomaly

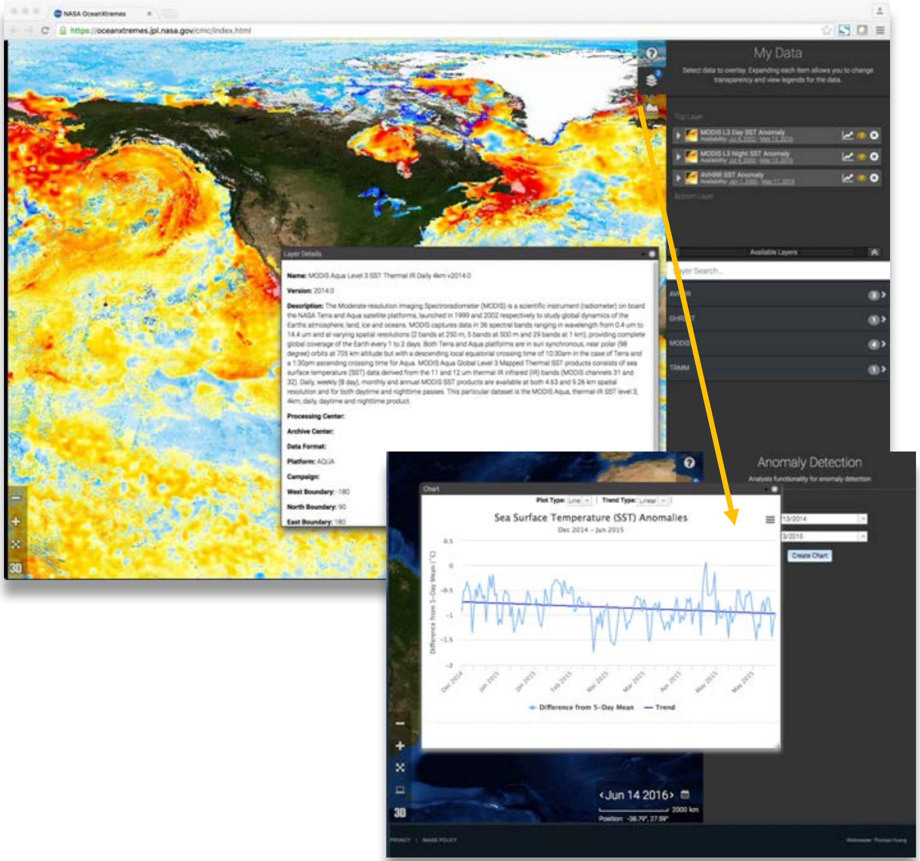

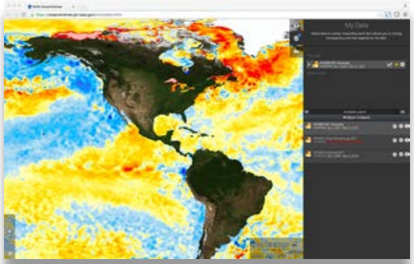

Aug 02, 2012

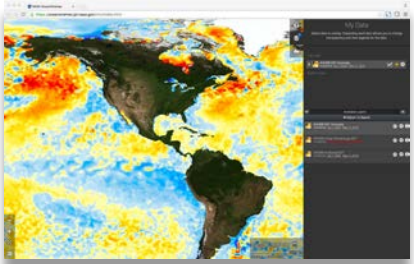

Aug 02, 2013

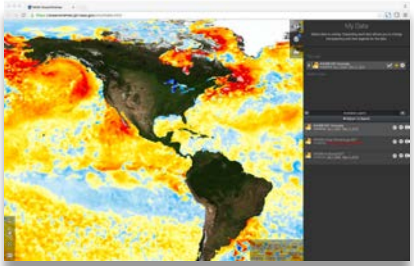

Aug 02, 2014

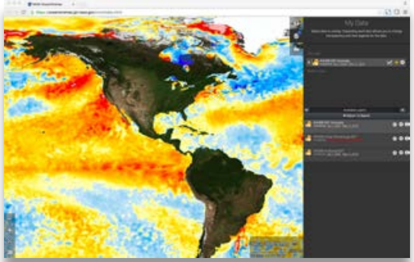

Aug 02, 2015

Matchup

Reconciliation of satellite and in-situ datasets

Distributed Oceanographic Matchup Service (DOMS)

- Funded by the NASA Advanced Information System Technology program
- 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

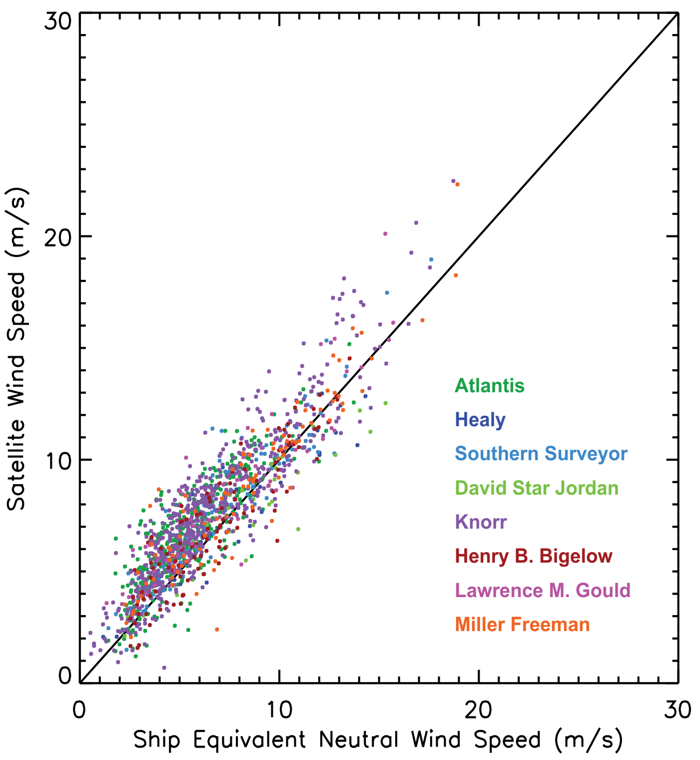

DOMS Data

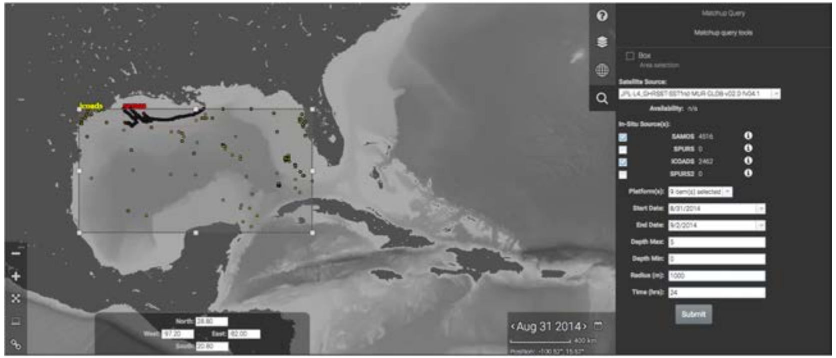

- **In-Situ**
  - **Shipboard Automated Meteorological and Oceanographic System (SAMOS)** initiative provides high-quality underway data from research vessels.
    - Hosted at Florida State University's Center for Ocean-Atmospheric Prediction Studies (COAPS), Tallahassee, Florida
    - Example OpenSearch query: `"http://doms.coaps.fsu.edu/ws/search/samos?startTime=2012-08-01T00:00:00Z&endTime=2013-10-31T23:59:59Z&bbox=-45,15,-30,30"`
  - **International Comprehensive Ocean-Atmosphere Data Set (ICOADS)** is a global ocean marine meteorological and surface ocean dataset.
    - Hosted at the National Center for Atmospheric Research (NCAR), Boulder, Colorado. Over 500 million measurements since year 1662.
    - Example OpenSearch query: `"http://rda-data.ucar.edu:8890/ws/search/icoads?startTime=2012-08-01T00:00:00Z&endTime=2013-10-31T23:59:59Z&bbox=-45,15,-30,30"`
  - **Salinity Processes in Upper Ocean Regional Study (SPURS)**
    - Hosted at the NASA Physical Oceanography Distributed Active Archive Center (PO.DAAC), Pasadena, California.
    - Example OpenSearch query: `"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"`
- **Satellite** – All managed by **NASA's Physical Oceanography Distributed Active Archive Center (PO.DAAC)**
  - Group for High Resolution Sea Surface Temperature (GHRSST) Level 4 Multiscale Ultrahigh Resolution (MUR)
  - Advanced Scatterometer (ASCAT) Level 2 coastal ocean surface wind vector
  - Soil Moisture Active Passive (SMAP) Level 2 Sea Surface Salinity (SSS)

National Aeronautics and Space Administration

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

DOMS Architecture

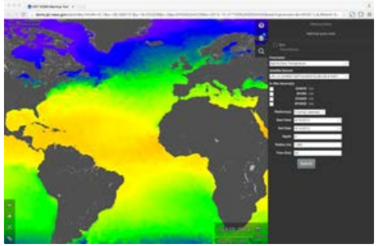

High resolution data visualization

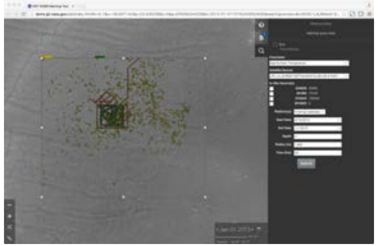

Initial match according to user initial selection

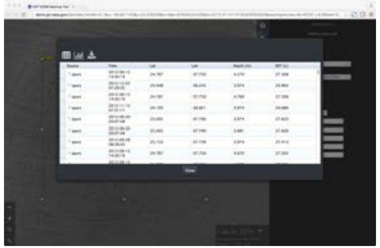

SST matchup with SPURS

Dynamic Match-up

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

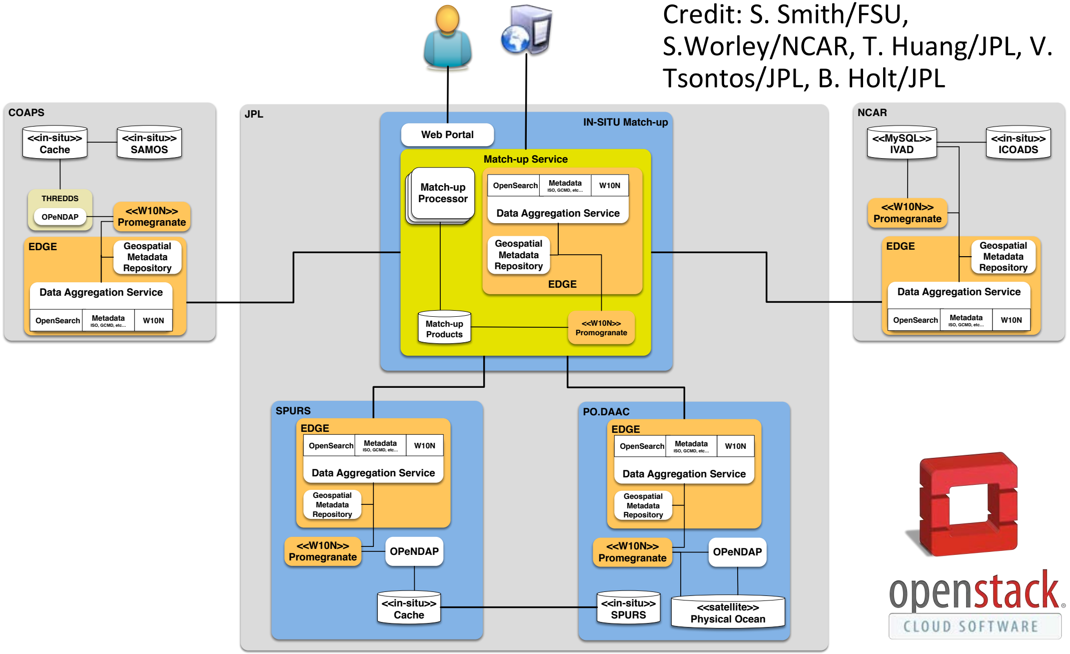


```
% curl -X GET
"https://doms.jpl.nasa.gov/nexus/match_spark?primary=JPL-L4_GHRSST-SSTfnd-MUR-GLOB-v02.0-fv04.1&matchup=spurs&startTime=2012-09-25T00:00:00Z&endTime=2012-09-30T23:59:59Z&b=-40,25,-45,30&platforms=1,2,3,4,5,6,7,8,9&depthMin=0.0&depthMax=5.0&tt=86400&rt=1000.0&parameter=sst"
```

Search and Discovery

Finding the right data and uncover related data and services

Search Relevancy

- Traditional keyword/key-phase search will not be adequate when dealing v petabyte-scale data
- What happen when a keyword/key-phase search returns thousands/millior hits? Which one should the user look at or download?
- **Search** – look for something you expect to exist
  - Information tagging
  - Indexed search technologies like Apache Solr or ElasticSearch
  - The solution is pretty straightforward
- **Discovery** – find something new, or in a new way
  - This is non-trivial
  - Traditional ontological method doesn't quite add up
  - The strength of semantic web is in inference
  - What happen when we have a lot of **`subClassOf`**, **`equivalentClassOf`**, **`sameAs`**?
  - How wide and deep should we go?
- **Relevancy**
  - It is domain-specific
  - It is personal
  - It is temporal
  - It is dynamic

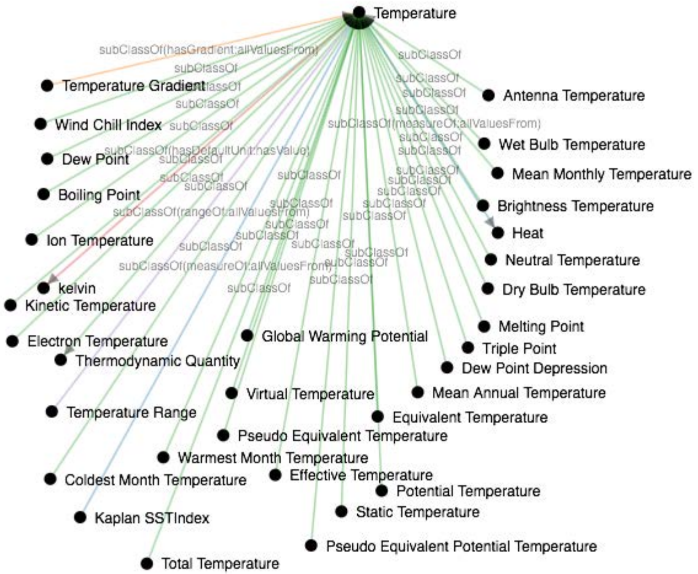

SWEET Ontologies' Temperature Concept
http://sweet.jpl.nasa.gov

Mining and Utilizing Dataset Relevancy from Oceanographic Datasets (MUDROD)

- Funded by the NASA Advanced Information System Technology program
- Analyze web logs to discover user knowledge (the connections between datasets and keyword)
- Construct knowledge base by combining semantics and profile analyzer
- Improve data discovery by
  - Better ranked results
  - Recommendation
  - Ontology navigation

✓ Web log processing
✓ Session reconstruction
✓ Vocabulary semantic relationship extraction
✓ Search ranking
✓ Recommendation

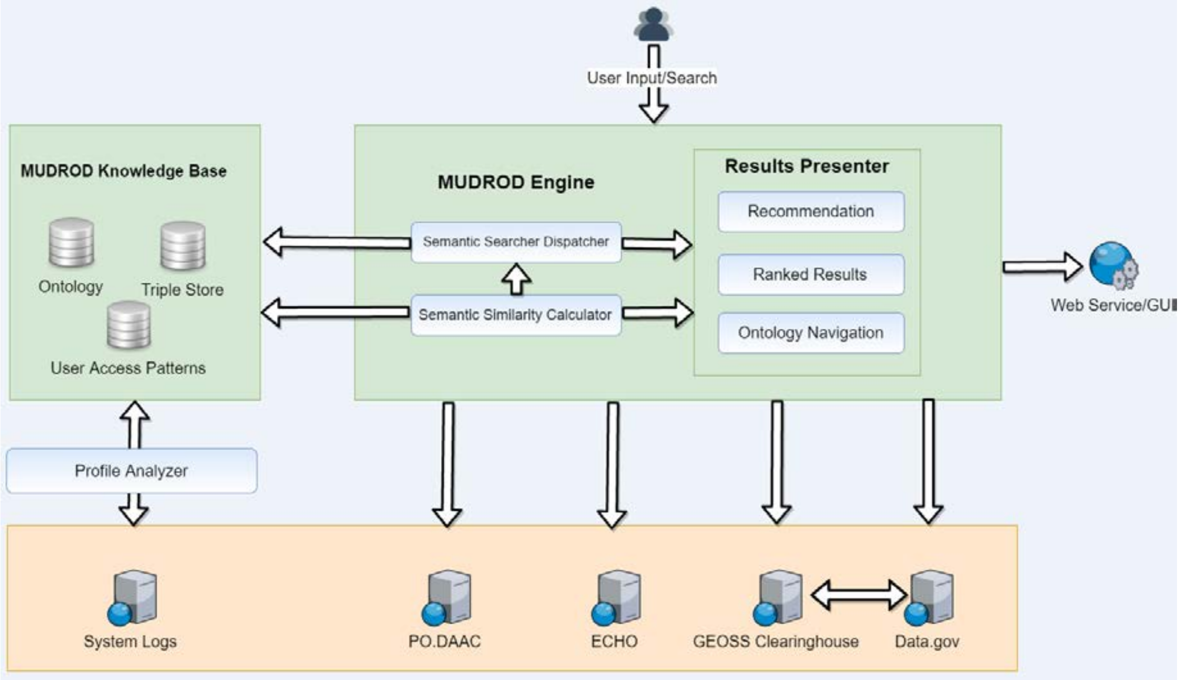

Key Takeaways

- OceanXtremes, DOMS, and MUDROD just finished their 1st year of development. Year 2 focuses on more performance updates and increase TRL level.
- More parallelism ≠ Faster performance
  - For global analytics, a lot of smaller tiles actually yields slower performance, because of scheduling, data transport, data queries, etc.
- Big Data ≠ 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

Further Information

- **NEXUS is available through open source – Apache License 2.0**
  - https://github.com/dataplumber/nexus

- **NEXUS at AGU 2016 Fall Meeting**
  - NASA Sea Level Change Portal – Boening, C. – Demonstration at NASA Booth
  - Session IN12A/B: Big Data Analytics – Huang, T., Lynnes, C., Vance, T., and Yang, C.
  - NEXUS-released Abstracts:
    - Gill, K., et.al, 2016: “Analysis of Sea Level Rise in Action”
    - Greguska, F, et.al, 2016: “Tackling the Four V’s with NEXUS”
    - Jacob, J., et.al, 2016: “Performance Comparison of Big Data Analytics With NEXUS and Giovanni”
    - Lynnes, C., et.al, 2016: “Benchmark Comparison of Cloud Analytics Methods Applied to Earth Observations”
    - Quach, N., et.al, 2016: “Sea Level Rise Data Discovery”
    - Wilson, B., et.al, 2016: “OceanXtremes: Scalable Anomaly Detection in Oceanographic Time-Series”

Special Thanks

NASA
Mike Little
Chris Lynnes
Kevin Murphy

FSU COAPS
Mark Bourassa
Jocelyn Elya
Shawn Smith
Adam Stallard

NCAR
Steve Worley
Ji Zaihua

GMU
Yongyao Jiang
Chaowei (Phil) Yang

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
David Moroni
Rob Toaz
Vardis Tsontos
Victor Zlotnicki

Questions, and more information

Thomas.Huang@jpl.nasa.gov