



**National Aeronautics and
Space Administration**

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

Lessons Learned in Creating Big Science Data Analysis Solution for the Cloud

Thomas Huang

thomas.huang@jpl.nasa.gov

Group Supervisor, Computer Science for Data Intensive Applications
Instrument Software and Science Data Systems

Strategic Lead for Interactive Analytics

National Space Technology Applications Program Office

[CL # 18-xxxx]

- **NASA has historically focused on systematic capture and stewardship of data for observational Systems**
- **With large amount of observational and modeling data, finding and downloading is becoming inefficient**
- **Reality with large amount of observational and modeling data**
 - Downloading to local machine is becoming inefficient
 - Search has gotten a lot faster. Too many matches.
 - Finding the relevant measurement has becoming a very time consuming process "*Which SST dataset I should use?*"
 - Analyze decades of regional measurement is labor-intensive and costly
- **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**
 - Capture 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


Processors are not Getting Faster

2004: First Pentium 4 processor with 3.0GHz clock speed

2018: Apple's MacBook Pro has clock speed of 2.7GHz

14 years later, not much has gain in raw processing power

42 Years of Microprocessor Trend Data




Original data up to the year 2010 collected and plotted by M. Horowitz, F. Labonte, O. Shacham, K. Olukotun, L. Hammond, and C. Batten
 New plot and data collected for 2010-2017 by K. Rupp


- **Mainly focus on archives and distributions**
- **With additional services**
 - Better searches – faceted, spatial, keyword, ranking, etc.
 - Data subsetting – home grown, OPeNDAP, etc.
 - Visualization – visual discovery, PO.DAAC's SOTO, NASA Worldview, etc.
- **Limitations**
 - Little to no interoperability between tools and services: metadata standard, keyword, spatial coverage (0-360 or -180..180), temporal representation, etc.
 - Making sure the most relevant measurements return first
 - Visualization is nice, but it doesn't provide enough information about the event/phenomenon captured in the image.
 - With large amount of observational data, data centers need to do more than just storing bits
 - “Is the red blob in the middle of Pacific normal this time of the year?”
 - “Any relevant news and publications relate to what I am looking at?”
 - “What other measurements, phenomena, news, publications relate to the period and location I am looking at?”
 - “I can see the observation from satellite, are there any relevant in situ data I can look at?”

Enabling Next Generation of Ocean Science Tools and Services


NASA Sea Level Change Portal



Oceanographic Anomaly Detection




PO.DAAC State Of The Ocean



Hydrological Basin Analysis




Jupyter Notebook - Interactive Workbench





Mobile Analysis




In Situ Data Analysis




Model Simulations




Model - Observation Comparison



Integrated Search and Discovery




Integrated Ocean Science Data Analytics Platform




- **Integrated Ocean Science Data Analytics Platform:** an environment for conducting a Ocean Science investigation
 - Enables the confluence of resources for that investigation
 - Tailored to the individual study area (physical ocean, sea level, etc.)
- Harmonizes data, tools and computational resources to permit the ocean research community to focus on the investigation
- Scale computational and data infrastructures
- Shift towards integrated data analytics
- Algorithms for identifying and extracting interesting features and patterns

Traditional Method for Analyze Satellite Measurements





- 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



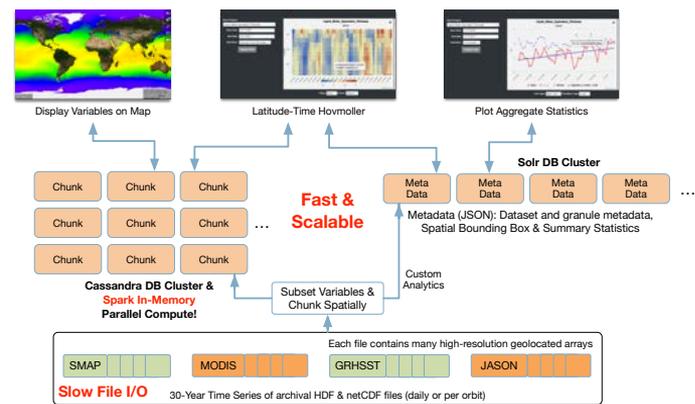
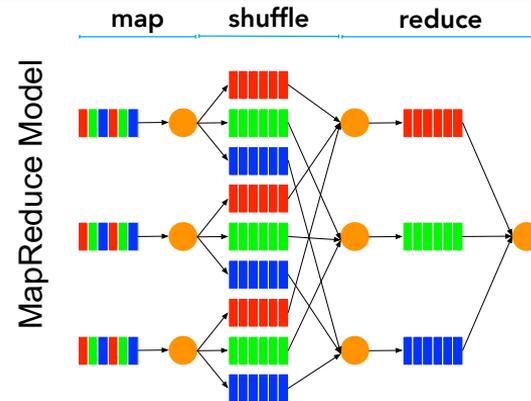
Observation

- 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



NEXUS: Scalable Data Analytic Solution


- MapReduce:** A programming model for expressing distributed computations on massive amount of data and an execution framework for large-scale data processing on clusters of commodity servers. - J. Lin and C. Dyer, *"Data-Intensive Text Processing with MapReduce"*
 - Map:** splits processing across cluster of machines in parallel, each is responsible for a record of data
 - Reduce:** combines the results from Map processes
- NEXUS** is 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 webservice




Two-Database Architecture

NEXUS' Pluggable Architecture for different Operation Needs

- **NEXUS supports public/private Cloud and local cluster deployments**
- **Several container-based deployment options**
 - Local on-premise cluster
 - Private Cloud
 - Amazon Web Service
- **Automate Data Ingestion with Image Generation**
 - Cluster based
 - Serverless (Amazon Lambda and Batch)
- **Data Store Options**
 - Apache Cassandra
 - ScyllaDB
 - Amazon Simple Storage Service (S3)
- **Resource Management Options**
 - Apache YARN
 - Apache MESOS
- **Analytic Engine Options**
 - Custom Apache Spark Cluster
 - Amazon Elastic MapReduce (EMR)
 - Amazon Athena (work-in-progress)





Sea Level Change - Data Analysis Tool

Visualizations | Hydrological Basins | Time Series | Deseason | Data Comparison | Scatter Plot |
 Latitude/Time Hovmöller | Etc.

NEXUS Performance: GIOVANNI vs. Custom Spark vs. AWS EMR

Dataset: MODIS AQUA Daily

Name: Aerosol Optical Depth 550 nm (Dark Target) (MYD08_D3v6)

File Count: 5106


Volume: 2.6GB

Time Coverage: July 4, 2002 – July 3, 2016

Giovanni: A web-based application for visualize, analyze, and access vast amounts of Earth science remote sensing data without having to download the data.

- Represents current state of data analysis technology, by processing one file at a time
- Backed by the popular NCO library. Highly optimized C/C++ library

AWS EMR: Amazon's provisioned MapReduce cluster **Giovanni: 20 min**
NEXUS: 1.7 sec



Area Averaged Time Series on AWS - Boulder

July 4, 2002 - July 3, 2016
 NEXUS Performance


Custom Spark vs. AWS EMR
 Ref. Speed - Giovanni: 1140.22 sec



Area Averaged Time Series on AWS - Colorado

July 4, 2002 - July 3, 2016
 NEXUS Performance


Custom Spark vs. AWS EMR
 Ref. Speed - Giovanni: 1150.6 sec



Area Averaged Time Series on AWS - Global


July 4, 2002 - July 3, 2016
 NEXUS Performance

Custom Spark vs. AWS EMR
 Ref. Speed - Giovanni: 1366.84 sec



Algorithm execution time. Excludes Giovanni's data scrubbing processing time

Analyze Ocean Anomaly – “The Blob”



- **Visualize** parameter
- **Compute** daily differences against climatology
- **Analyze** time series area averaged differences
- **Replay** the anomaly and visualize with other measurements
- **Document** the anomaly
- **Publish** the anomaly





Figure from Cavole, L. M., et al. (2016). "Biological Impacts of the 2013–2015 Warm-Water Anomaly in the Northeast Pacific: Winners, Losers, and the Future." *Oceanography* 29.

Enable Science without File Download



```
# 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=mesos,16,32&ds=AVHRR_OI_L4_GHRSSST_NCEI&minLat=45&minLon=-150&maxLat=60&maxLon=-120&startTime=1220227200&endTime=1443657600

It took: 2.9428272247314453 sec
```

Using IDL with NEXUS

```

IDL> spawn, 'curl'
" https://oceanworks.jpl.nasa.gov/timeSeriesSpark?spark=mesos,16,32&ds=AVHRR_OI_L4_GHRSS
T_NCEI&minLat=45&minLon=-150&maxLat=60&maxLon=-120&startTime=2008-09-
01T00:00:00Z&endTime=2015-10-01T23:59:59Z" -o json_dump.txt'
  
```


% Total	% Received	% Xferd	Average Speed	Time	Time	Time	Current
			Dload Upload	Total	Spent	Left	Speed
0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0
0	353k	0	0	0	0	0	0
72	353k	72	256k	0	52705	0	0:00:02
100	353k	100	353k	0	69303	0	0:00:05

```

IDL>
IDL> result = JSON_PARSE('json_dump.txt', /toarray, /tostruct)
IDL> help, result
** Structure <1a2749c8>, 3 tags, length=62320, data length=62320, refs=1:
  STATS      STRING      '!NULL'
  META       STRUCT      -> <Anonymous> Array[1]
  DATA      STRUCT      -> <Anonymous> Array[778]
  
```

```


IDL>
IDL> plot(result.data.time, result.data.mean, title='GHRSSST L4 AVHRR_OI SST. Sep 2008
- Oct 2015. US West Coast Blob Area')
PLOT <29457>
  
```



Credit: Ed Armstrong
Jun. 05, 2018



Hurricane Katrina Study




Powered by NEXUS

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

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 °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.



Hurricane Katrina
 TRMM
 overlay SST
 Anomaly

Multi-Variable Analysis


- Public accessible RESTful analytic APIs where computation is next to the data
- NEXUS as the analytic engine infused and managed by the DAACs on the Cloud
- Researchers can perform multi-variable analysis using any web-enabled devices without having to download files
- Reduce unnecessary egress charges
- An architecture to enable next generation of scientific applications



Architecture for Distributed Data System and Analysis



Committee on Earth Observation Satellites (CEOS)

- A common service registry
- Putting value-added services next to the data to eliminate unnecessary data movement
- Avoid data replication
- Public accessible RESTful analytic APIs where computation is next to the data
- Analytic engine infused and managed by the data centers perhaps on the Cloud
- Researchers can perform multi-variable analysis using any web-enabled devices without having to download files
- Reduce unnecessary data movement and egress charges
- An architecture to enable next generation of scientific applications




In Situ to Satellite Matchup

- Distributed Oceanographic Matchup Service (DOMS)
- 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.
- Plugin architecture for in situ data source using EDGE, an open source implementation of Open Search




NASA AIST OceanWorks: Ocean Sciences Platform

- **OceanWorks** is to establish an **Integrated Data Analytic Center** at the NASA Physical Oceanography Distributed Active Archive Center (PO.DAAC) for Big Ocean Science
- Focuses on technology integration, advancement and maturity
- Collaboration between JPL, FSU, NCAR, and GMU
- Bringing together PO.DAAC-related big data technologies
 - Anomaly detection and ocean science
 - Big data analytic platform
 - Distributed in-situ to satellite matchup
 - Search relevancy and discovery – linking datasets, services, and anomalies through recommendations
 - Metadata translation and services aggregation
 - Fast data subsetting
 - Virtualized Quality Screening Service



Serverless Ingestion Workflow Architecture

- The cluster ingestion workflow requires several EC2 instances and attached EBS storage. Amazon bills active EC2 and storage even there is no data to be processed
- Cost-saving serverless architecture using AWS S3 Bucket triggering AWS Lambda job to create AWS Batch jobs



Analytic Storage



Blocked
Storage
(EBS)

\$0.045/GB-month
\$47,186/PB-month

V.S.



Object
Storage
(S3)

\$0.021/GB-month
\$22,020/PB-month


NEXUS supports 2 types of cloud storage

- Blocked storage (e.g. Amazon EBS), Attach to computing node. Generally faster
 - Cassandra and ScyllaDB
- Object storage (e.g. Amazon S3), Independent storage service. Highly scalable

For a data center, not all data need to be served on fast storage, Object storage provides a better, scalable alternative

Architecture for Custom Algorithms and Cluster Sharing

- Developed independently, all the major services in OceanWorks require Apache Spark cluster.
- If OceanWorks simply deploy these services to Amazon, it will require dedicated Apache Spark cluster for each.
- Too many clusters and very costly, especially for high memory machine instances.
- Looking at the Amazon's EMR model. It is designed to be a job execution solution, and the jobs could from different applications.
- Apache Livy provides a RESTful interface to Apache Spark cluster.
- We built upon the Livy solution and extended the NEXUS Command Line interface to enable "push" of custom code for execution on OceanWorks' remote Spark cluster.
- Scientists can use the Jupyter environment to design their analytic algorithms that will be executed in the OceanWorks' Spark cluster.




Push Python/Scala Code to Spark Cluster

- Users can now push their own custom map-reduce algorithms in Python/scala to OceanWorks for execution in the cloud.
- For our Python users, this is part of the NEXUSCLI library. It is Jupyter friendly and standalone python applications
- Enables services to share Apache Spark computing cluster.
- **New REST endpoints:**
 - https://oceanworks.jpl.nasa.gov/run_str
 - Run custom code in a (multi-line) string.
 - https://oceanworks.jpl.nasa.gov/run_file
 - Run custom code in a local file.
 - These new endpoints are not public.
 - For trusted (authenticated) users.


New NEXUS CLI Python library methods:

```
from nexuscli import nexuscli_ow
nexuscli_ow.set_target('https://oceanworks.jpl.nasa.gov')
code = '[x*x for x in range(11)]'
ans = nexuscli_ow.run_str(code)
fname = 'custom_code.py'
ans = nexuscli_ow.run_file(fname)
```


Support for Hydrology



Retrieval of a single river time series



Retrieval of time series from 9 rivers




Time series coordination between TRMM and river

- Simulated hydrology data in preparation for SWOT hydrology
- **River data: ~3.6 billion data points.** 3-hour sample rate. Consists of measurements from ~600,000 rivers
- **TRMM data: 17 years, .25deg, 1.5 billion data points**
- Sub-second retrieval of river measurements
- On-the-fly computation of time series and generate coordination plot

Tackling Information Discovery

- Adjust search result according how user search and retrieval
- Use Machine Learning approach to adjust search ranking by taking a number of features into consideration – version, processing level, release date, all-time popularity, monthly-popularity, and user popularity
- Semantically mind dataset metadata to identify relationship
- **Relevancy** is
 - Domain-specific
 - Personal
 - Temporal
 - Dynamic



Dynamic Ranking and Recommendation

Search results for 'RSCAT_LEVEL_2B_OWV_CLIM_12_V1'.

Long Name: RapidScan Level 2B Climate Ocean Wind Vectors in 12.5km Footprints

Topic: Surface Winds

Category: Earth Science

Variable: Surface Winds

Term: Ocean Winds

Release Date: 06/05/16

Abstract: This dataset contains the RapidScan Level 2B 12.5km Version 1.0 Climate quality ocean surface wind vectors. The Level 2B wind vectors are based on a 12.5 km Wind Vector Cell (WVC) grid and processed using the using the "full aperture" normalized radar cross-section (NRCS), a.k.a. Sigma-0) from the L1B dataset. RapidScan is a Ka-band dual beam coastal scanning synthetic-aperture radar (SAR) system with the same hardware and functionality of CHIRP/SAR, with exception of the antenna sub-system and digital receiver to the International Space Station (ISS) Calipso receiver, which is where RapidScan is mounted. The NASA mission is officially referred to as ISS Synthetic Aperture Radar (ISAR). RapidScan is not a sun-synchronous orbit, and flies at roughly half the altitude with a low inclination angle that restricts data coverage to the tropics and mid-latitude regions, the extent of latitudinal coverage depends from approximately 36 degrees North to 36 degrees South. Furthermore, there is no consistent local time of day retrieval. This dataset is provided in a netCDF-3 file format that follows the netCDF-4 API and made available via FTP and OPENDAP. For data access, please click on the "Data Access" tab above. This climate quality data set offers from the nominal "full" L2B dataset as follows: 1) it uses full antenna footprint measurements (20 km) without subsampling by range (1/3 km and 2) the absolute calibration has been completed for the low SNR signal to noise ratio (SNR) mode data sets. LowSNR1: 14 August 2015 to 18 September 2015, LowSNR2: 6 October 2015 to 7 February 2016. The above enhancements allow this dataset to provide consistent calibration across all SNR states. As a result of the frequent and unpredictable transition between high and low SNR states, this dataset will only be updated when new low SNR data transitions occur. The plan is updating new data to the dataset time series above the calibration to be better controlled over longer time periods, thus resulting in a final data product that is more suitable for climate research. Low SNR periods and other key quality control (QC) issues are tracked and kept up-to-date here: ftp://podas.gsfc.nasa.gov/Data/data/dds/catalog/?v=system.csw. If you have any questions, please visit our user forums: https://podas.gsfc.nasa.gov/forum/

Processing Level: 2

DOI: 10.5067/REK12.L2C11

TemporalRepeat: 12 Hour

TemporalResolution: 3 Day

TemporalRepeatMin: 1.5 Hour

Order: RapidScan

Project: ISS_RapidScan

Format: -


DataLatency: 168

Related datasets:

- RSCAT_LEVEL_2B_OWV_COMP_12_V1_2
- RSCAT_LEVEL_2B_OWV_COMP_12_V1_1
- RSCAT_LEVEL_2B_OWV_COMP_12
- OSCAT_LEVEL_2B_COMP_12
- SEAWINDS_LEVEL_2B_COMP_12
- RSCAT_L1A_12KM_V1_2
- RSCAT_L1B_V1_2
- OS2_OSCAT_LEVEL_2B_OWV_COMP_12_2
- RSCAT_L1A_28KM_V1_2
- OSCAT_LEVEL_2B_V2


Related searches:

- surface wind (1)
- wind (0.99)
- wind speed (0.93)
- quikscat (0.88)
- vector (0.76)
- wind data (0.79)
- topex position wind (0.75)
- quikscat wind (0.74)



Container-based deployment using Docker




- Dockerized all service components
- Register with DockerHub
- Integration with Continuous Integration (CI) solutions (e.g. Jenkins or Atlassian Bamboo)
- Same container deployment process for on-premise hardware or commercial Cloud environment




Example Container Deployment of our Open Source Analytic Engine, NEXUS

Deployment Automation

- Cloud Deployment is nontrivial
- Infrastructure Definition
 - Various machine instances
 - Storage and buckets
- Software Deployment.. manually
 - Build
 - Package
 - Install
 - Configure
 - Shell login (security issues)
- Best Practice: Deployment Automation
 - Script Infrastructure Definition (e.g. Amazon CloudFormation)
 - Container-based Deployment (e.g. Amazon ECS and DockerHub)




AIST Managed Cloud Environment



ESDIS NGAP

Free and Open Open Source Software (FOSS)

- October 2017, the OceanWorks project released all of its source code to Apache Software Foundation and established the **Science Data Analytics Platform (SDAP)** in the **Apache Incubator**
- Technology sharing through Free and Open Source Software (FOSS)
- Why? Further technology evolution that is restricted by projects / missions
- It is more than GitHub
 - Quarterly reporting
 - Reports are open for community review by over 6000 committers
 - SDAP has a group of appointed international Mentors: Jörn Rottmann, Raphael Bircher, and Suneel Marthi
- OceanWorks is now being developed in the open
 - For local cluster and cloud computing platform
 - Fully containerized using Docker (multiple containers)
 - Infrastructure orchestration using Amazon CloudFormation
 - Analyzing satellite and model data
 - In situ data analysis and colocation with satellite measurements
 - Fast data subsetting
 - Data services integration architecture
 - OpenSearch and dynamic metadata translation
 - Mining of user interactions and data to enable discovery and recommendations
 - Streamline deployment through container technology



<http://sdap.apache.org>



Notable Public Engagements

- Host hands-on cloud analytics workshops using Amazon Web Services (AWS)
- Invited to speak at the **Space Studies Board of The National Academy of Sciences**
- Invited to present to the **NASA Advisory Council's Ad-Hoc Big Data Task Force (BDTF)**
- Invited to present to the **JPL Deputy Lab Director and Chief Technologist**
- Invited to present to **CNES Chief Technologist and Delegations**



In Summary

- Traditional method for scientific research (search, download, local number crunching) is unable to keep up
- Let's think beyond archive and file downloads
- Connected information enables discovery
- Community developed solution through open sourcing
- Investment in data and computational sciences
- Data Centers might want to be in the business of Enabling Science!
- OceanWorks infusion 2018 – 2019
 - Turnkey AMCE deployment
 - Solution for large job management
 - Integration with Amazon Athena
 - Integration with Pangeo
 - Watch for changes to the NASA's Sea Level Change Portal
 - Even faster analysis capabilities
 - More variety of measurements – satellites, in situ, and models
 - Event more relevant recommendations
 - NASA's Physical Oceanography Distributed Active Archive Center (PO.DAAC)
 - More than just pretty pictures. SOTO will have new analytic capabilities.
- Coming Soon: 2018 Wiley Book on **Big Earth Data Analytics in Earth, Atmospheric and Ocean Sciences**



**National Aeronautics and
Space Administration**

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

Thomas Huang

Thomas.huang@jpl.nasa.gov

Jet Propulsion Laboratory
California Institute of Technology