Adapting Critical Space Domain Awareness Applications to Run at Large Scale on a Cloud-Based Data Analytics Platform

Chang Zhang, Ann L. Chervenak, Alex Gonring, Mark Mendiola, Jeffrey Won, Scott Bergonzi, Derek Chen, Vincent Kong, Eltefaat Shokri

March 1, 2021

Motivation

- As the number of space objects proliferates, there is a growing need to perform automated, high-fidelity orbit analysis on a large number of these objects

- **Transitioning legacy high-fidelity orbit analysis and other Space Domain Awareness (SDA) applications to cloud architectures empowers scalable and modern usage**

- We present a general approach for cloud-based, high-fidelity, scalable orbit analysis
  - *We leverage high-fidelity orbit propagation as an initial use case*

- Our approach:
  - ***Containerize legacy orbit propagation software*** *(Aerospace Corporation's TRACE software)*
  - ***Deploy multiple instances*** *of the containerized TRACE application* ***on a cloud-based data analytics platform*** *(Aerospace Corporation's Data Exploitation (DEX) platform)*
  - ***Download state vector data*** *from LeoLabs and Numerica* ***from the Unified Data Library (UDL)***
  - ***Perform orbit propagation*** *on these state vectors in parallel using TRACE instances deployed on DEX*
  - ***Visualize the results***
    - **Show propagated trajectories** for these objects (using NASA WorldWind)
    - **Provide aggregate statistics** for propagation performance

- This approach shows promise for integrating and scaling other legacy SDA applications

State Vectors Processed by TRACE Analytic

- **State vector data is from LeoLabs and Numerica**
  - *LeoLabs operates multiple radar stations around the globe and tracks thousands of objects in Low Earth Orbit*
  - *Numerica operates a worldwide telescope network; fuses third-party optical, radar, and passive-RF measurements*
  - *Both companies publish observations, state vectors and uncertainty estimations for each tracked object to the Unified Data Library*
- **Unified Data Library (UDL)**
  - *UDL is a repository for Space Domain Awareness data from many sources*
  - *Being developed by BlueStaq for Air Force Research Laboratory (AFRL) and Space and Missile Systems Center (SMC)*
- **Our experiments query the UDL to retrieve state vectors based on:**
  - *Specified source (e.g., LeoLabs, Numerica)*
  - *NORAD ID value(s)*
  - *Epoch range*

The DEX Data Analytics Platform

- The cloud-based Data EXploitation (DEX) platform was developed under SMC Enterprise Ground
- DEX is a **Reference Architecture** and **prototype implementation** for a data analytics platform
  - *Enables the exploitation of available and relevant data sources to identify abnormal behaviors in real-time*
- The DEX platform provides a software layer between data analytics applications and the underlying cloud infrastructure to manage the applications
- Supports **integration of new and legacy analytics applications**
- Extracts information by running configurable, **general-purpose and domain-specific analytics**
- Supports delivery of **actionable, real-time information** for satellite operations
- Provides **intuitive user interfaces** that allow users to **configure and deploy analytics workflows** easily

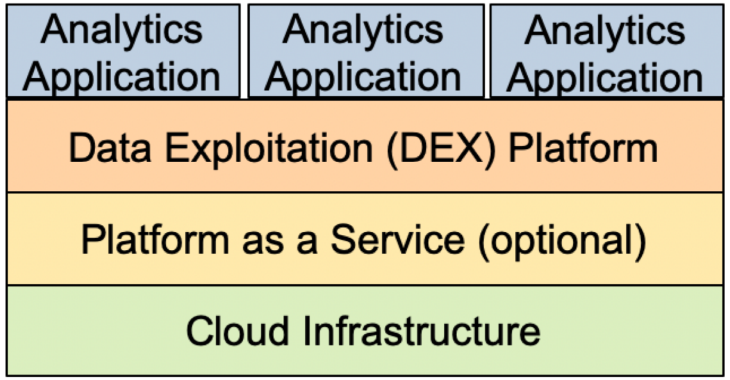

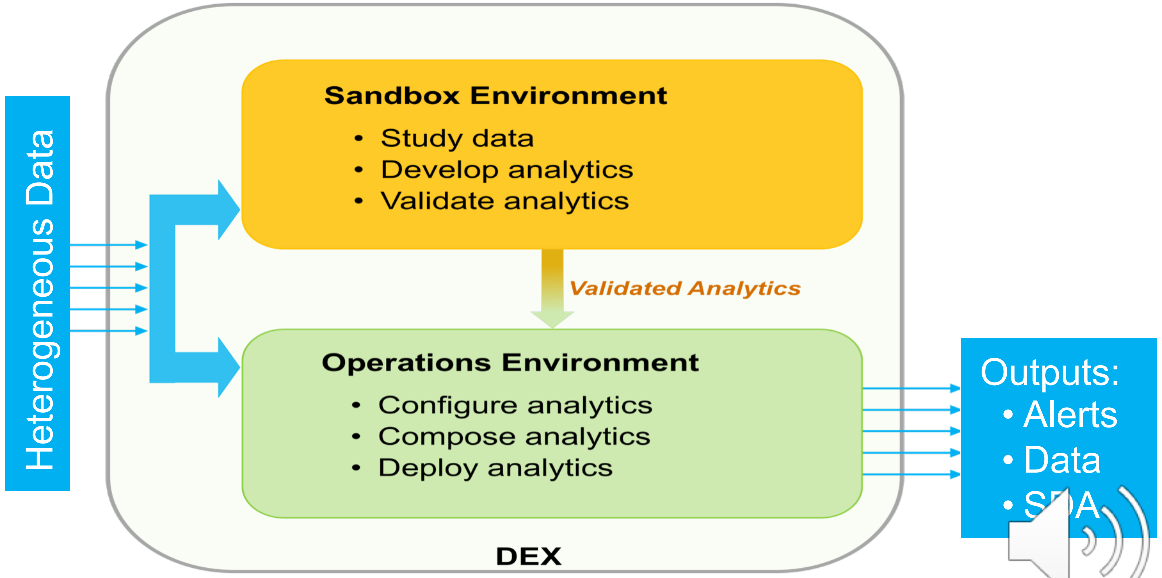

System Integration

- Two analytic types deployed on DEX cloud-based analytics platform:
  - *__UDL Data Download Analytic__ that queries LeoLabs and Numerica state vector data from UDL*
  - *Multiple instances of __TRACE orbit propagation analytic__ that generates propagated trajectories for downloaded state vectors*
- Analytics utilize Apache Kafka Data Streaming for communication within DEX
- Visualization with NASA WorldWind Visualization

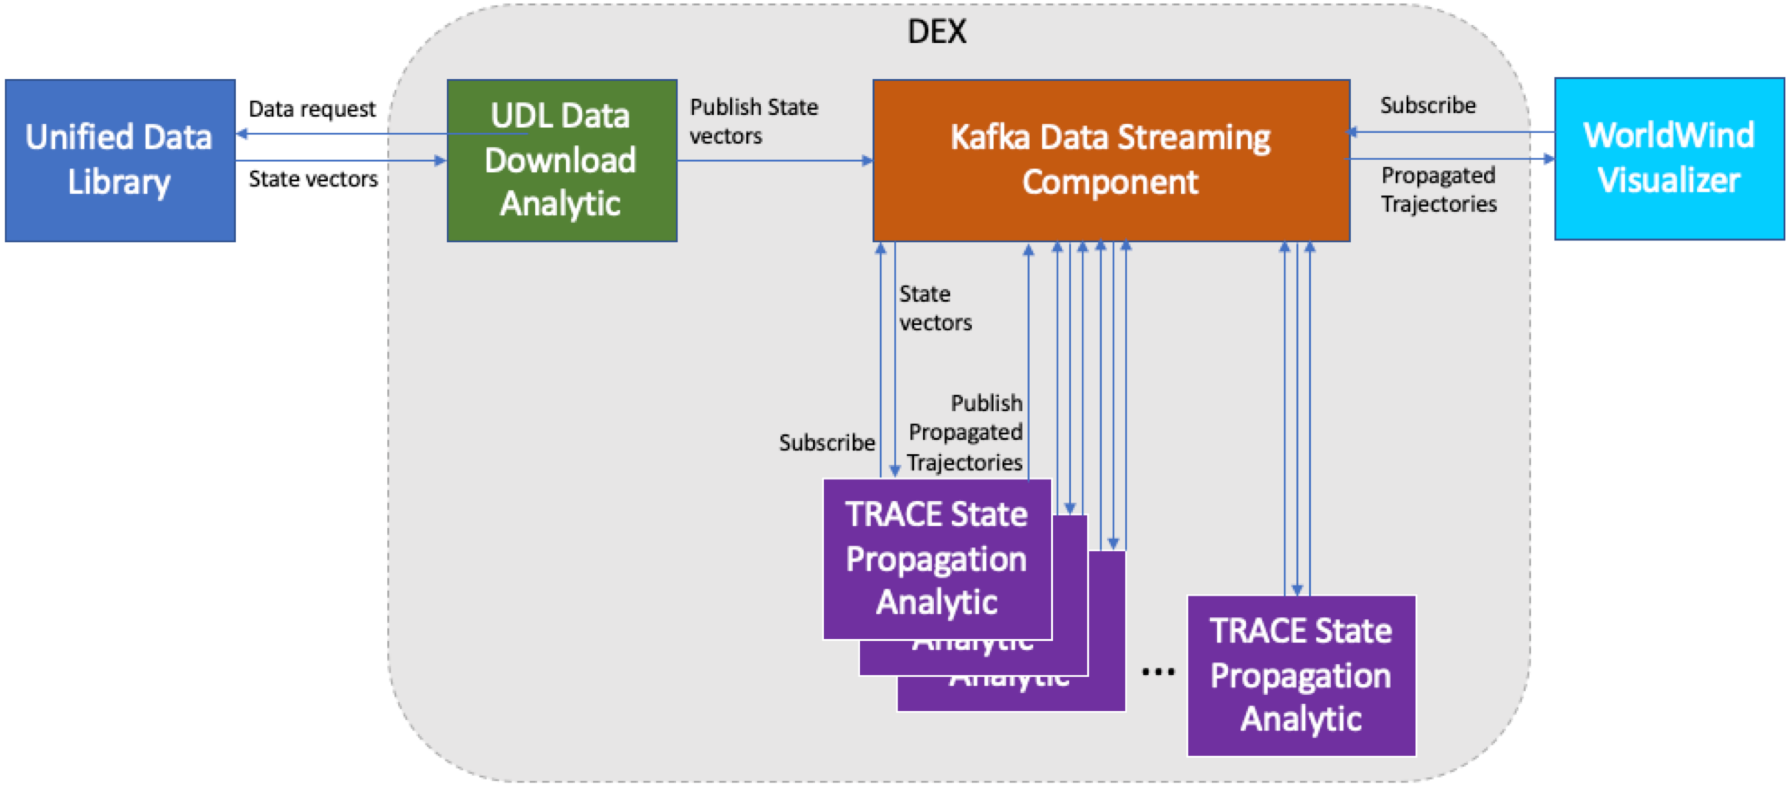

TRACE

- **TRACE is the Aerospace Corporation's trajectory analysis and orbit determination program**
  - *Continuously developed since the 1960s*
- TRACE has been used regularly to **support a variety of high-fidelity analyses**
  - *Used as benchmark for independent verification and validation of commercial astrodynamics tools and orbit propagators*
  - *Used extensively in support of various U.S. government space programs*
- TRACE capabilities include:
  - ***Orbit propagation and ephemeris generation***
  - ***Measurement data generation***
  - ***Measurement error modeling***
  - ***Orbit determination***
  - ***Covariance analysis***
- TRACE provides a high level of detail and configurability available for the user
  - *Ability to model orbital trajectories with a high-order numerical integrator or a selection of analytic propagators*
  - *Numerical integrator allows user to configure many detailed force models to characterize the dynamic environment*
  - *Output files provide detailed ephemerides and intermediate results of the user's choosing, including states in multiple coordinate frames, partial derivatives, and special event print options for times of eclipse and nodal crossing*
- **TRACE is a complex tool** due to the many options available to the user and numerous high-fidelity models
  - *One goal of containerization is to* ***provide a simpler interface to a subset of TRACE functionality***

Containerization of the Legacy TRACE Application

- **Containerization is a virtualization technique used in cloud computing**
  - *Most common virtualization techniques: Virtual Machines and Containers*
  - *Virtualization allows sharing of physical resources (processors, memory, networks, etc.) of a single physical machine by multiple applications*
  - ***Containerization:** An application is packaged in a software container with all its software dependencies*
  - *Containers are lightweight, portable and can be deployed quickly on a cloud*
- The containerized TRACE analytic only **exposes the specific TRACE functionality related to orbit propagation**
  - *Expose only a portion of the overall TRACE capabilities to **provide focused and simplified use of a complex engineering tool***
  - *Requires only a simplified set of input parameters to reduce the complexity for the user*
- Packaged in a Docker container with all dependent libraries and databases

Required Input Parameters
Stop Epoch
Time Step
Gravity Model Degree and Order
State Input Coordinate Frame
Group_ID
Partition_Assignment

TRACE Support Databases and Files
Earth Orientation Parameters
Leap Seconds Table
Canned Planetary Ephemerides
Canned EGM2008 Gravity Model
Planetary Geomagnetic Amplitudes: Ap
Solar Flux Indices: F10.7

Containerized TRACE Application

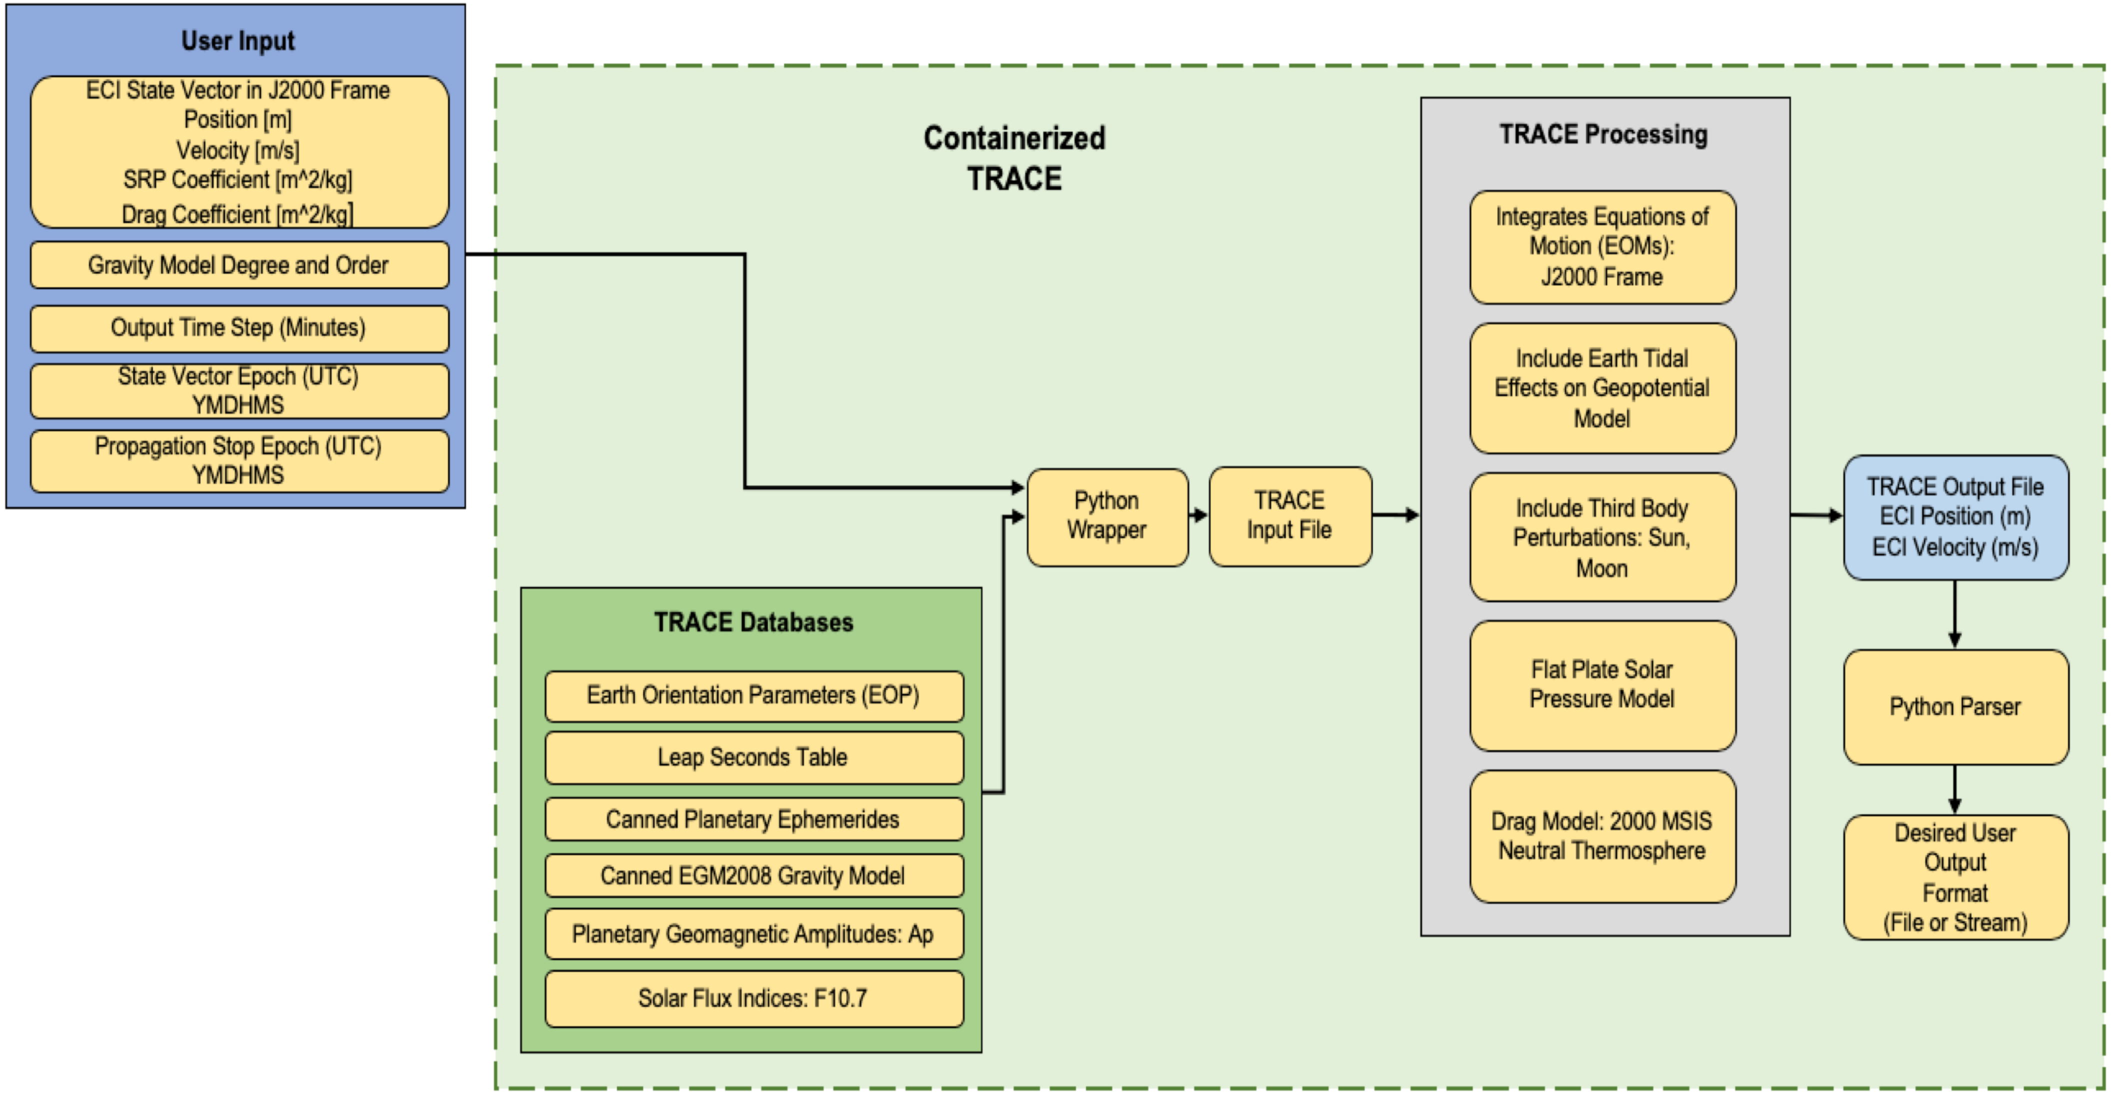

UDL Data Download Analytic

- For a specified source, NORAD ID value(s) and epoch range, this analytic **queries the Unified Data Library (UDL) to retrieve state vectors**

- The containerized analytic does the following:
  - *Uses input parameters to* ***create UDL state vector query***
  - ***Issues the query to the UDL***
  - ***Receives state vector results of query***
  - ***Publishes state vectors to Apache Kafka topic***
    - Topic has configurable number of partitions

Required Input Parameters	Description
token	Required to access UDL
source	String specifying the data source (e.g., LeoLabs or Numerica)
norad_id	A list of North American Aerospace Defense Catalog Numbers, also known as Satellite Catalog Numbers
epoch	Start epoch string for the UDL query
end_epoch	End epoch string for the UDL query
max_results	Specify maximum number of results for UDL query
total_partitions	Specifies number of Kafka partitions to be used for load balancing in DEX

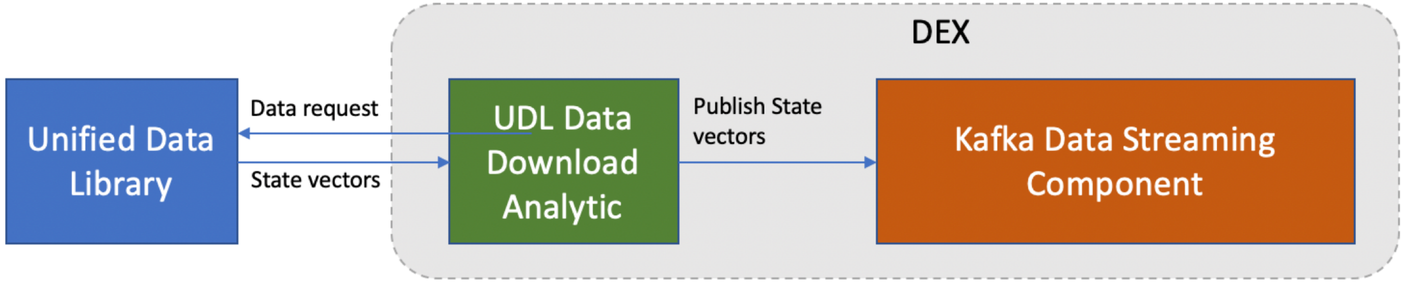

Experiment 1: Processing 10 LeoLabs objects in parallel on cloud

- Deployed 10 instances of TRACE analytic, one instance of UDL Data Download analytic on DEX Platform
- NASA WorldWind visualization shows propagated orbit trajectories for these objects
- Presented in AMOS 2020 paper

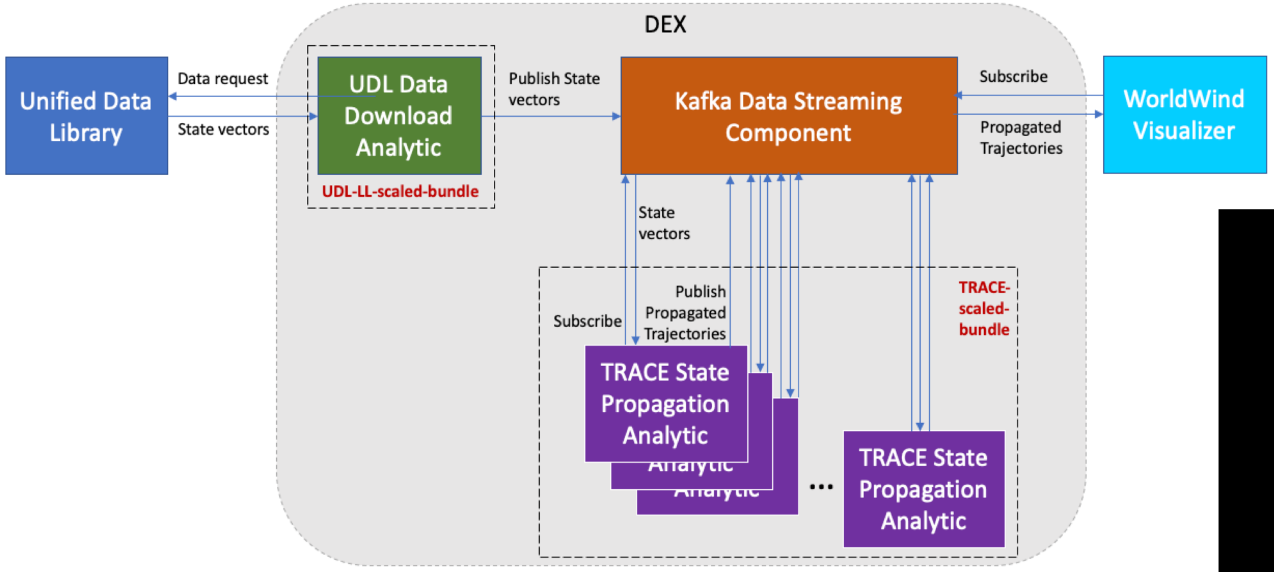

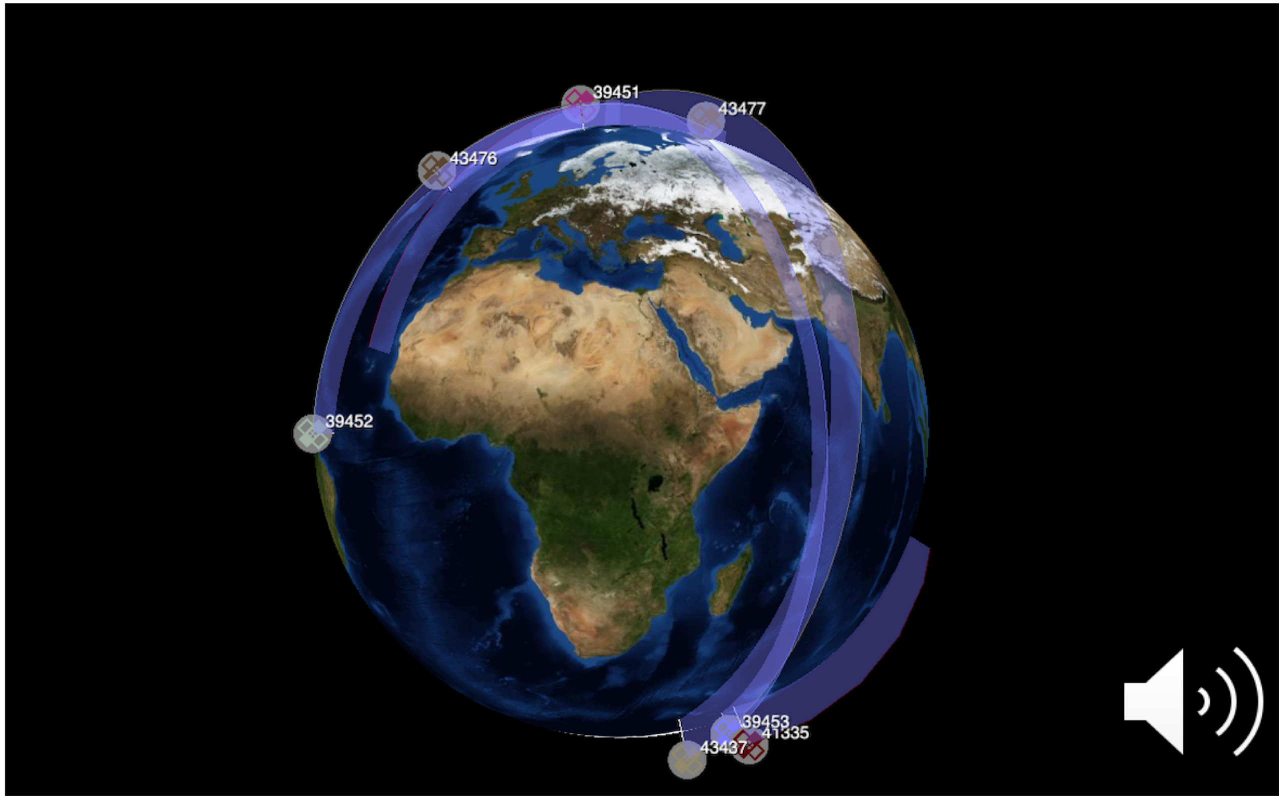

Experiment 2: Processing an Increasing Number of Numerica Objects in Parallel with 20 Instances of TRACE

- Fixed number of TRACE analytic instances deployed on a DEX cluster (20)
- Increasing number of state vectors downloaded from UDL (100 to 2000)
- Propagated each state vector for 24 hours
  - *Output one state per minute, 1441 per state vector*
- Scales well to 2000 objects: Linear increase as number of state vectors increases

State Vectors from UDL	Time Elapsed for Processing (seconds)	# of TRACE instances deployed	Number of state vectors propagated per second
100	18.97	20	5.27
200	41.61	20	4.81
400	85.29	20	4.69
600	129.98	20	4.62
800	172.72	20	4.63
1000	219.3	20	4.56
1200	260.82	20	4.60
1400	304.05	20	4.60
1600	347.11	20	4.61
1800	389.28	20	4.62
2000	439.05	20	4.56

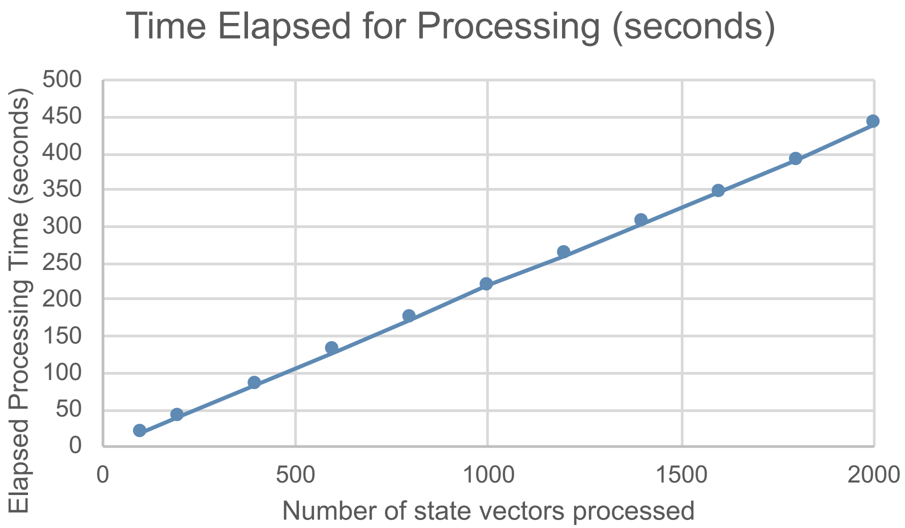

Experiment 3: Processing 1000 Numerica Objects with an Increasing Number of TRACE Instances

- Fixed number of objects downloaded from UDL (1000)
- Propagated each state vector for 24 hours
  - *Output one state per minute, 1441 per state vector*
- Increasing number of instances of TRACE deployed on cluster (1 to 50 instances)
- Diminishing benefits after 25 instances (due to cluster configuration)

State Vectors from UDL	Time Elapsed for Processing (seconds)	Number of TRACE Instances Deployed	Number of state vectors propagated per second
1000	3694.87	1	0.27
1000	524.27	10	1.91
1000	309.71	20	3.23
1000	219.13	25	4.56
1000	207.93	30	4.81
1000	208.07	50	4.81

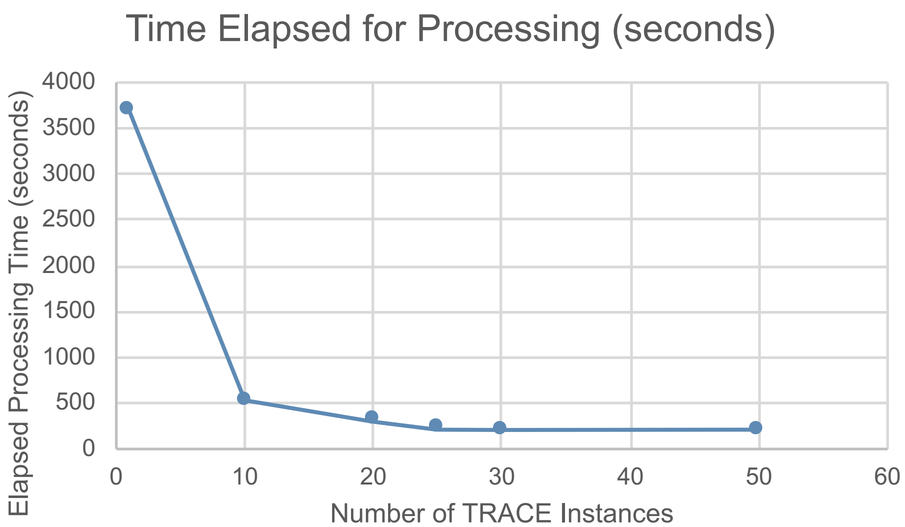

Summary and Future Work

- Goal of this work is to adapt critical SDA applications to run at large scale on a cloud-based data analytics platform
  - *To support the urgent need for analysis of large number of objects in congested, contested space*
- Demonstrated that the DEX cloud-based data analytics platform can scale to support the analysis of thousands of space objects
- Showed the ability to containerize, deploy and scale legacy SDA applications on the cloud
  - *Example: TRACE trajectory analysis and orbit determination program*
  - *Containerized the TRACE orbit propagation capability*
- Future work:
  - *Continue to scale number of objects analyzed & number of TRACE instances*
  - *Create additional analytics that containerize and expose other TRACE capabilities*
  - *Containerize and integrate additional Space Domain Awareness applications*