

Dragster 2.0: An operations-ready framework for neutral density assimilation

Rachel Stutz, Jeff Steward, Connor Johnstone, Tyler O’Connell, Junk Wilson, John Noto
Orion Space Solutions, Louisville, CO

Marcin Pilinski
*Laboratory for Atmospheric and Space Physics (LASP), University of Colorado at Boulder,
Boulder, CO*

David Vallado
COMSPOC Corp., Exton, PA

Shaylah Mutschler
Space Environment Technologies (SET), Pacific Palisades, CA

ABSTRACT

We present a quantitative comparison of three methods to modeling the neutral density for low earth orbit (LEO) altitudes including the novel data assimilative approach represented by Dragster. The first approach uses an empirical neutral density model output driven by solar and lower altitude forcing features but uninformed by physics. This method is used by the Naval Research Laboratory (NRL) in the Mass Spectrometer and Incoherent Scatter Radar Extended model (NRLMSIS-00). The second method uses physics to drive thermospheric and ionospheric conditions from a starting state. This method is used by the High Altitude Observatory (HAO) at the National Center for Atmospheric Research (NCAR) in the Thermosphere-Ionosphere-Electrodynamics General Circulation Model (TIE-GCM) 2.0. Finally, we present Dragster 2.0, a unique commercial application to find maximum likelihood, minimum variance estimates of thermospheric neutral density along with quantified uncertainty. It uses satellite ephemerides in an ensemble Kalman filter based architecture with the possibility of running over 1000 members of the background density model and performing assimilation of the derived effective density observations in each member. Dragster 2.0 features a newly rearchitected framework appropriate for operational use based on Rust and REST API microservices. The ensemble Kalman filter assimilates effective density derived from satellite state vectors with a background neutral density model. This allows calculation of mean and spread statistics from the ensemble distribution of any variables that are part of the model state. These include both corrections on neutral density values and inputs to the background model such as solar forcing terms (F10.7 and Kp/Ap). The final model states and uncertainties are stored in a single PostgreSQL database, allowing for easy querying for both application of Dragster corrections to neutral density models and validation of Dragster’s performance across many test cases. Presented here are results from an empirical (NRLMSIS) atmospheric model as a background to Dragster, calculating the ensemble spread for the model state. These distributions provide quantified confidence intervals for neutral density values in the LEO region as well as solar forcing terms during the time of model runs. Dragster outputs a grid of multiplicative factors that can be applied as corrections to a background neutral density grid. The improved density estimates and their uncertainties can be used as a background for satellite propagation, leading to improved space domain awareness (SDA). Including these calculated confidence intervals for neutral density in propagation algorithms also has the potential to improve the accuracy of conjunction assessments for space traffic management (STM). The resulting solar forcing terms and their uncertainties can be used in other assimilative forecast schemes to improve modeling and understanding the space environment. In addition, the ability to use the cloud for high performance computing (HPC) is beneficial for research and development as well as profiling in operations. We present our experience with Parallel Works, a vendor who specializes in cloud-based HPC workflows. This platform has enabled our scientists and engineers to scale up and test the deployment of our TIE-GCM container solution on the cloud. This allows us to test the impact of a large number of TIE-GCM ensemble members (1000+) on Dragster performance to determine the optimum cost versus performance. Also presented are preliminary results of visualizing Dragster-updated neutral densities from NRLMSIS as well as the difference between stand-alone and Dragster-updated NRLMSIS density grids in the CesiumJS framework along with concurrent satellite positions and orbits. The results of the three sources of neutral density (NRLMSIS, TIE-GCM, and Dragster) are compared for SDA and STM use cases. The outputs of the three orbital propagation calculations are visualized and the results of multiple ensemble members with perturbed solar forcing parameters and background density values are compared. This study informs cost and value trade off decision makers considering each of the three neutral density methods for their SDA and STM requirements. This study should also inform government policy makers considering how to provide public services required for safety of flight while still preserving a competitive commercial market for commercial STM applications that foster innovation. The final goal is a comprehensive system that can help scientists and operational decision makers understand the “what now” (current neutral density conditions and satellite orbits), “what next” (error-quantified forecast orbits influenced by density corrections), and “what if” (scenario exploration such as investigating solar storms) of the space domain.

1. INTRODUCTION

Satellites in low-earth orbit (LEO) experience a variety of forces in the dynamic environment of the thermosphere. The forcing from solar storms can lead to orders of magnitude differences of neutral density compared with quiet conditions. The drag experienced as a result of these thermospheric changes can make accurately predicting spacecraft positions challenging. Space Domain Awareness (SDA) applications include conjunction assessments to avoid potentially catastrophic collisions. In order to propagate a satellite position forward in time, in addition to orbital dynamics terms, accurate specification of the neutral density (and neutral wind) in the orbital region is needed for accurate drag prediction [8]. Other applications where improved satellite forecast can benefit include object tracking and maneuver detection.

Uncertainty bounds on satellite forecasts are extremely important in conjunction assessments and other SDA use cases, but quantified uncertainty is hard to come by in space weather modeling and forecasting [13]. Variations in solar and geomagnetic forcing and the resulting variations in atmospheric parameters such as density are highly non-Gaussian, making it difficult to derive accurate statistical representations of the atmospheric state. The use of ensembles in space weather modeling has the potential to address this gap in understanding. Each ensemble member provides a version of the model state, capturing the effect of input variations as the model runs and providing at the end a distribution of model states from which statistical parameters can be calculated.

2. MODEL AND APPLICATION OVERVIEWS

We compare three methods of obtaining neutral density in the thermosphere:

1. The Mass Spectrometer and Incoherent Scatter Radar Extended model (NRLMSISE-00 or NRLMSIS), from the Naval Research Laboratory (NRL) and the University of Colorado Space Weather Technology, Research and Education Center (SWx TREC).
2. The Thermosphere-Ionosphere-Electrodynamics General Circulation Model (TIE-GCM) 2.0, from the High Altitude Observatory (HAO) at the National Center for Atmospheric Research (NCAR).
3. The Dragster 2.0 application, from Orion Space Solutions, an Arcfield Company (OSS).

An overview of each model or application used in this study is included below, with a more comprehensive description of Dragster provided in Section 3.

2.1 NRLMSIS

The NRL Mass Spectrometer and Incoherent Scatter Radar model (NRLMSIS) is an empirical model of the atmosphere [14]. It provides species and total density and temperature estimates based on temporal and solar conditions. For this study, the `pymis` Python package was used to run the model, using indices downloaded from CelesTrak [1] [10] [11].

2.2 TIE-GCM

The Thermosphere-Ionosphere-Electrodynamics General Circulation Model (TIE-GCM) is a physics-based model of the I-T system [17]. Given space weather conditions parameterized by solar and geomagnetic indices, TIE-GCM propagates the state of the space environment forward in time through simulations of the physical interactions between resident particles. For this study, TIE-GCM 2.0 was run at 5-degree resolution with driving indices from GFZ [3] [11].

2.3 Dragster

Dragster is an assimilative thermospheric neutral density application that incorporates satellite ephemeris data in LEO to update a background density model. It was first developed at ASTRA, LLC. (now Orion Space Solutions / OSS) along with government and university partners [15] [16]. The application has since undergone significant refactoring efforts to improve its scalability and operational capability. Version 2.0 features the code body rewritten from IDL to Rust and a modular approach with containerized microservices. Dragster uses an ensemble Kalman filter to assimilate observations, resulting in statistical parameters on the model state that can provide uncertainty on density and driver outputs. For this study, Dragster 2.0 was run with NRLMSIS as its density background.

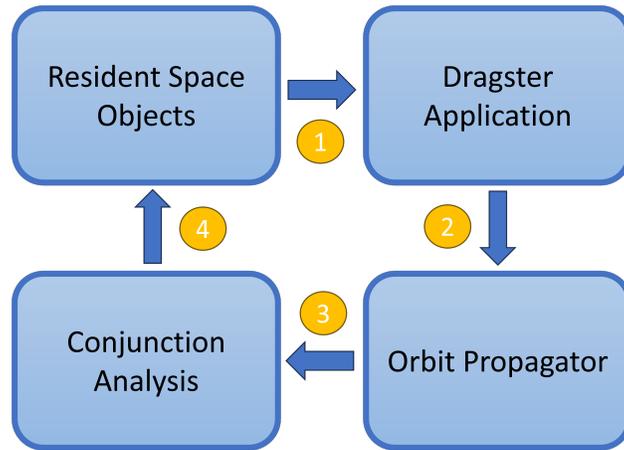


Fig. 1: Dragster as an SDA tool. Process is 1) assimilate observations 2) access density results 3) update orbit predictions 4) situational awareness.

3. DRAGSTER 2.0 ARCHITECTURE

Dragster provides an up-to-date picture of the thermosphere through advanced data assimilation methods. A high-level depiction of Dragster used in an SDA flow is shown in Fig. 1. Through the use of an ensemble Kalman Filter (EnKF), Dragster provides a distribution of ensemble members at each time step of the model run. Statistical parameters drawn from the ensemble distribution can provide users with calculated uncertainty estimates on the entire model state, which includes both solar and geomagnetic drivers and total density throughout the thermosphere. A visualization of the total density output on the globe after running Dragster on the NRLMSIS model is shown in Fig. 3. Ongoing upgrades will allow the Dragster 2.0 microservices to be parallelized and run in high-performance computing (HPC) environments, allowing for computationally expensive runs with large ensembles of physics-based background models and physics-based forecasting capabilities.

The key features of Dragster 2.0 are described in the following subsections.

3.1 Containerized Microservice Architecture

Dragster is composed of a series of modular components which are containerized and communicate to each other via API requests. These components are either Rust Actix-Web microservices, simpler Rust compiled libraries, Python Fast-API microservices, PostgreSQL databases, or Apache Airflow Directed Acyclic Graphs (DAGs). Each service with its type and a brief description of its function is shown in Fig. 2. These services are intended to allow for streamlined deployment of Dragster for operational (continuous, real-time) use. OSS has an operational deployment of Dragster that was set up after refactoring of the Dragster code was completed, and developments continue to be pushed to the operational deployment. The containerized microservice architecture makes it possible to update and build one service at a time and then swap out the container image used in a given deployment stage. Extract-Transform-Load (ETL) pipelines manage data ingestion and application runs. The Data Ingestion Pipeline is an Airflow DAG used as an ETL pipeline to automatically retrieve newly-published spacecraft ephemeris data packets and pass them to the ingestion service for storage in the observations database. When deployed, another pipeline kicks off runs of Dragster on a regular cadence by making a request to the Dragster Main Service, which orchestrates the remainder of the run.

Since the background model is its own microservice that is called during a Dragster run via an API request (Background Density Service in Fig. 2), the specific model can be switched out as desired. Dragster 2.0 was originally built to run with NRLMSISE-00 as the background model, but another thermospheric density model such as TIE-GCM, Jacchia-Bowman 2008 (JB-08), or the Whole Atmosphere Model (WAM) can be used instead. The inputs to the background model are included in Dragster's state that is updated with the EnKF, so updates and uncertainty estimates for these inputs as well as the total density are available after a Dragster run. A visualization of how Dragster data assimilation updates NRLMSIS total density outputs is shown in Fig. 5.

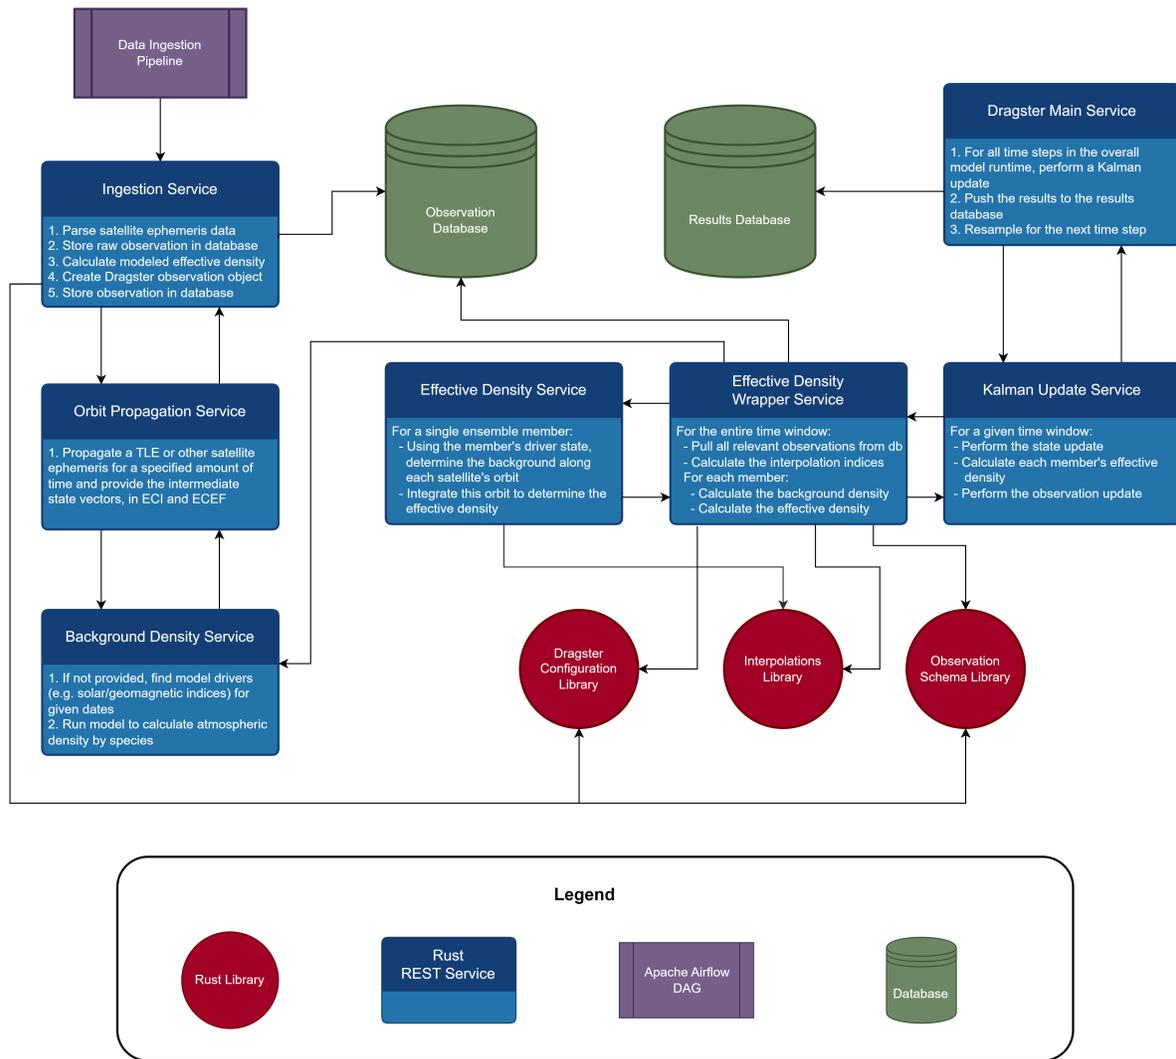


Fig. 2: Dragster 2.0 architecture with microservices.

3.2 Data Assimilation

Dragster incorporates observations at each time step to bring its background density model closer to reality [9]. Effective total density observations derived from satellite ephemeris data in LEO are assimilated and used as observations in the EnKF. Historically, Dragster has used two-line element sets (TLEs) as these observations, since they are the most widely publicly available form of spacecraft positional data. In version 2.0, Dragster is now able to ingest a series of state vectors, which tend to have higher fidelity and accuracy than TLEs. Ephemeris data in state vector format can be retrieved from places such as the Unified Data Library [6].

3.3 Ensemble Kalman Filter

Dragster uses an ensemble Kalman Filter (EnKF) to update its background model with observations. The use of an EnKF for Dragster is introduced in [16]. The ensemble members represent a distribution of model states, from which mean and spread statistics and thus quantified uncertainty can be calculated at the end of a Dragster run.

3.4 Uncertainty

Dragster 2.0 provides the ensemble mean, standard deviation, skewness, and kurtosis for each point in the output density grid as fields in the results database. The true distribution can only be approximated with these four parameters, but they are intended to be a lightweight representation of the ensemble (which can be quite large with hundreds or

thousands of members), allowing the user to calculate rough positional uncertainty values when used in propagation. Quantified uncertainty on a spacecraft's predicted position is very useful for SDA use cases such as conjunction assessment and avoidance. Uncertainty characterization derived from the full ensemble will be available in an upcoming Dragster update, as described in Section 3.7.

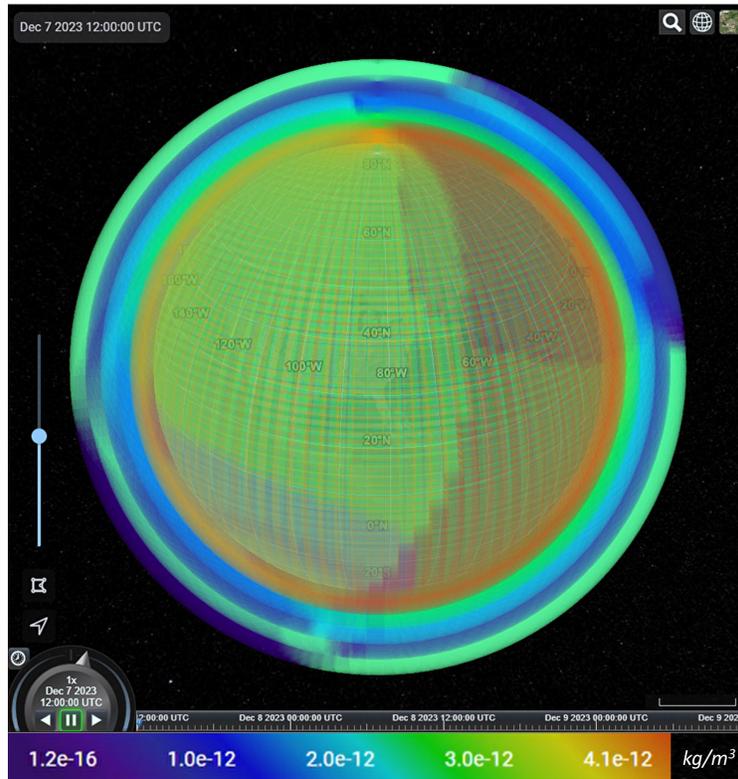


Fig. 3: Dragster total density outputs on 8 December 2023, visualized using CesiumJS.

3.5 HPC Integration

Dragster is bringing the ability to harness High-Performance Computing (HPC) for more powerful Dragster runs in ongoing development work. This makes it easier to run physics-based models (such as TIE-GCM) as backgrounds for Dragster with large numbers of ensemble members, which can be extremely computationally intensive. Running such expensive tasks in a traditional cloud or on-premise deployment can be limiting for operational models and applications. HPC takes advantage of additional computational resources to run complicated, expensive models in a more time-efficient manner, and Dragster is also being designed to use Slurm clusters to burst out to the powerful compute nodes only when necessary, keeping costs reasonable. This architecture is being designed to be vendor-agnostic through Parallel Works, who provide HPC control planes for provisioning compute clusters [4].

3.6 Physics-Based Forecasting

Dragster can integrate physics-based forecasting of density outputs. This relies on the use of physics-based models such as TIE-GCM being updated by Dragster and then using the model to forecast into the future after assimilating past and current observations. HPC assists greatly in this capability due to the computationally demanding nature of physics-based models; thus, forecasting is also ongoing alongside the HPC integration. Neutral winds are also being integrated into Dragster alongside this development.

3.7 Full Ensemble Outputs

Dragster updates also include saving the output state of the entire ensemble rather than just a few statistical parameters meant to adjust a Gaussian distribution. The total density outputs from the full ensemble can be used for a more accurate representation of the spread, especially in SDA use cases. Propagating a spacecraft using the full ensemble

densities provides a more accurate uncertainty estimate than one obtained using just a few parameters, since this approach takes into account the entire spread of the distribution and results in a “cloud” of potential propagated positions. This could be used to characterize an uncertainty ellipsoid around a spacecraft with some quantified confidence, as depicted in Fig. 4. If the ellipsoid represented a 95% confidence interval, for example, then a well-tuned ensemble should find a validating measured position for the spacecraft to be within the ellipsoid 95% of the time.

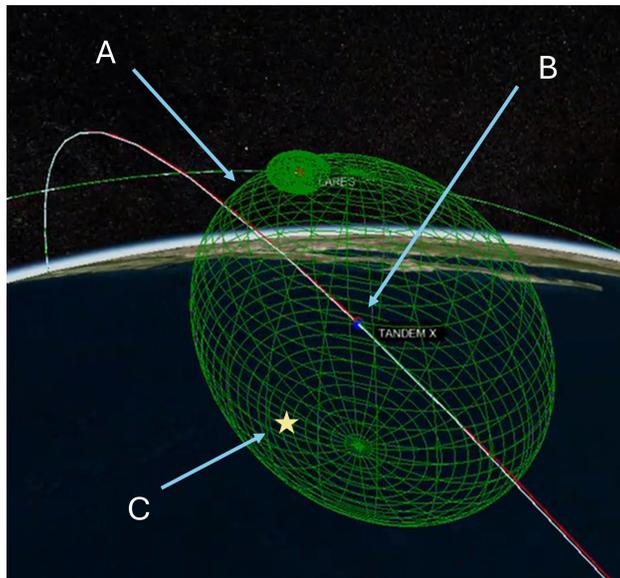


Fig. 4: Sample ensemble positional uncertainty ellipsoid (A) surrounding a satellite propagated using Dragster 2.0 outputs (B), with sample validating measured position (C) within the ellipsoid.

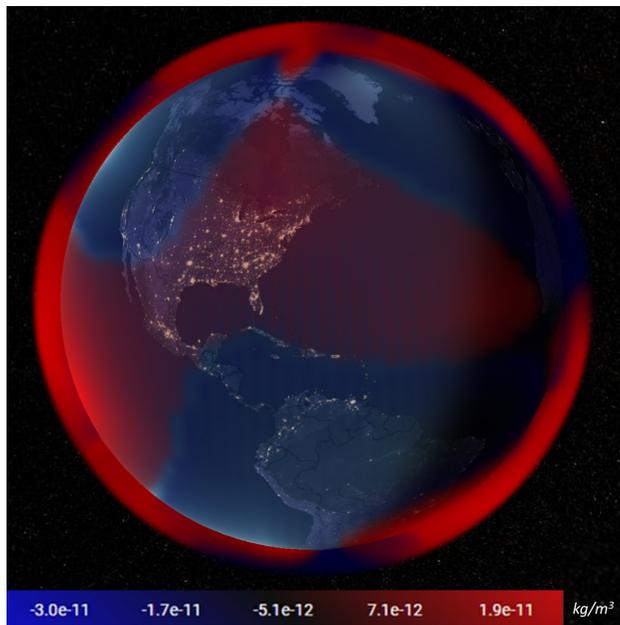


Fig. 5: Difference between Dragster-updated NRLMSIS and stand-alone NRLMSIS total density outputs on 8 December 2023. Red indicates that Dragster output densities were higher, whereas blue indicates that stand-alone NRLMSIS output densities were higher.

4. ORBIT PROPAGATION TESTING

The performance of each of the thermospheric density outputs from the models and application mentioned in Section 2 can be indicated by how accurately a propagator using that density output can predict where a spacecraft will be at some later point in its orbit. Assuming that the spacecraft does not maneuver during the period of propagation or accounting for maneuvers, the modeled drag for the spacecraft is very important, as drag induced by the neutral density in the orbital regime is the most significant contributor to error in trajectory prediction for LEO spacecraft [8]. Therefore, propagating forward from a measured starting position, the most accurate density output should result in the spacecraft's propagated position being closest to its measured position.

For this comparison, a LEO spacecraft was propagated using density outputs from each of the thermospheric density models (NRLMSIS, TIE-GCM) and the Dragster 2.0 application between a measured starting and ending position over a period of approximately one day. The chosen spacecraft orbits in the 600 km altitude range, placing it within the altitude range of all density models.

The output time of interest for the comparison was during the Gannon solar storm (10-12 May 2024). MSIS and TIE-GCM ingest the daily $F_{10.7}$ solar index and the 3-hour Kp and/or ap geomagnetic indices as proxies for geomagnetic conditions during the model run. These indices were particularly heightened during this storm, with ap reaching 400 on 10 May in the afternoon UTC and Kp reaching 9.0 twice on 11 May [12]. This allowed us to test performance during extreme geomagnetic storm conditions. The three-hour geomagnetic Kp indices for two three-day windows, one before and one during the Gannon storm, are shown in Fig. 6.

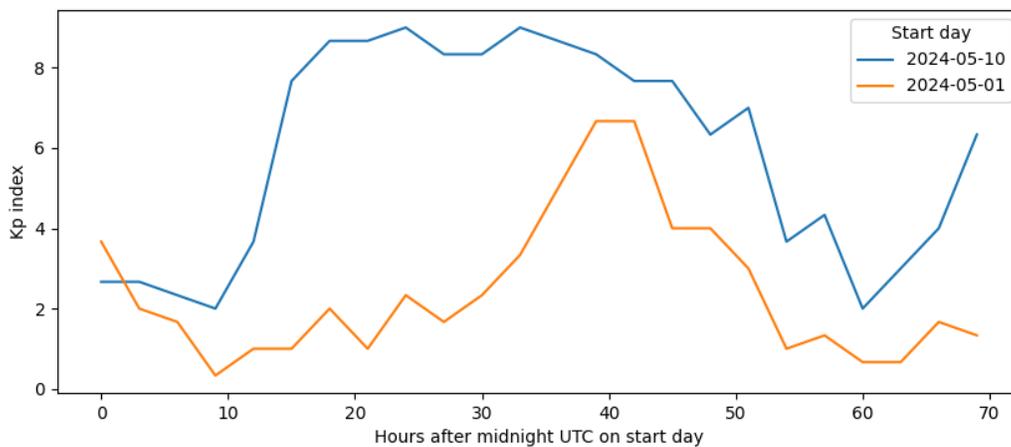


Fig. 6: Kp indices during the two propagation intervals in May 2024 [3].

Dragster assimilated TLEs from up to 96 non-maneuvering LEO satellites during the testing period. The effective density that each satellite experienced during its orbit was derived from the ingested state vector and TLE measurements to update the background total density output. For this test, the background density model was NRLMSISE-00, accessed via the `pymis` library and running in a Docker container. The satellite chosen for propagation testing was not in the set of satellites used for data assimilation.

The output densities in each case were integrated using Orion Space Solutions' in-house orbit propagation library, which allows for a custom density model to be passed in for modeling drag on the spacecraft being propagated. In addition to specifying the density, the parameters shown in Table 1 were provided to the propagator. The starting state vector for the spacecraft and the measured ending state vector used for error calculations were both derived from TLEs downloaded from `space-track.org` [5] using Julia's `SatelliteToolbox` library [7].

From this starting state, the satellite was propagated using each density model (or Dragster-updated model) for drag specification until the time of another measured position approximately one day later. The difference between the measured position and the propagated position for each case was computed, and the results are shown in Section 5.

Table 1: Satellite parameters used for orbit propagation tests.

Satellite	COSMOS 921
NORAD ID	10095
Start Position	[-4.160730e6, 5.616426e6, 41.792289]
Start Velocity	[-1472.525004, -1111.236275, 7311.563879]
Start Time (UTC)	2024-05-10 16:54:06.403
End Time (UTC)	2024-05-11 17:00:28.450
Mass (<i>kg</i>)	5000
Area (<i>m</i> ²)	19.635

5. RESULTS

The difference between the test satellite’s propagated ending positions and measured ending position for each density background are shown in Table 2. The reported error is the shortest distance in space between the measured ending position and each propagated ending position (i.e. the norm of the difference between the two state vectors). The test was performed near solar maximum and during a geomagnetic storm, as indicated by the $F_{10.7}$ and Kp/ap indices. The propagation times for these spacecraft were limited by TLE availability and the satellite’s altitude remaining within the altitude bounds of NRLMSIS and TIE-GCM (TIE-GCM is more restrictive, with its upper altitude bound in the 600-700 km range). Many more such tests should be performed to provide more comprehensive comparisons of these thermospheric density models.

Table 2: Propagation error after 1 day (10-11 May 2024) for the COSMOS 921 spacecraft.

Model	Error (km)
NRLMSIS	2.269
TIE-GCM	3.417
Dragster	2.058

Visualizations of space weather models can be useful in understanding the global variations in atmospheric density and other space weather products as well as the difference between outputs from various models. A few examples of such visualizations, using CesiumJS as the visualization for both an Earth observation and orbital digital twin, are shown [2]. Fig. 5 provides a comparison of stand-alone and Dragster-updated NRLMSIS global total density outputs at a given time. Fig. 7 provides a look at variation in spacecraft position after propagating with different density/perturbation models. Finally, Fig. 8 displays a spacecraft flying through different density models that could be used as drag backgrounds for propagation. The difference in densities at various geospatial locations can clearly be seen with the datasets side-by-side in the split screen view, demonstrating the scientific potential of the advanced visualizations available in such an integrated Earth and satellite digital twin.

6. CONCLUSIONS

Our results indicate that data assimilation and ensemble approaches can be effective at providing updated neutral density estimates in the thermosphere within a commercial operations-ready architecture. Importantly, this does not showcase a statistically significant number of tests of the performance of these models. Space environmental conditions as well as the quality of satellite ephemeris data assimilated could affect model performance. The Dragster application will be further tested and validated against existing models with more propagation tests under multiple different temporal, solar, and geomagnetic conditions, as well as using additional satellite observations and background models.

Further improvements to Dragster 2.0 will be validated as they are incorporated into the production application. For example, the physics-based TIE-GCM has some advantages in density specification over the empirical NRLMSIS, especially during storm times. Using TIE-GCM as a background for Dragster would provide the additional advantage

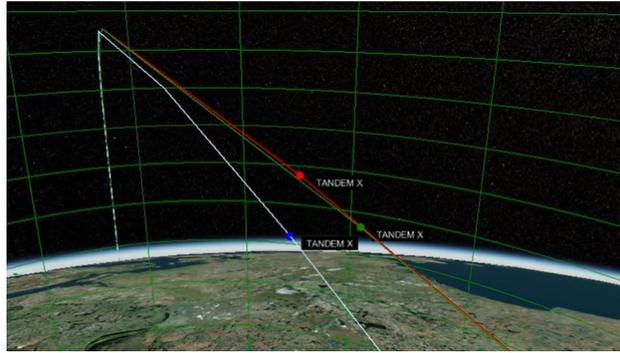


Fig. 7: TANDEM-X satellite propagated with different background density models visualized in CesiumJS: SGP4 (blue dot / white line), NRLMSIS (green dot/line), and Dragster-updated NRLMSIS (red dot/line).

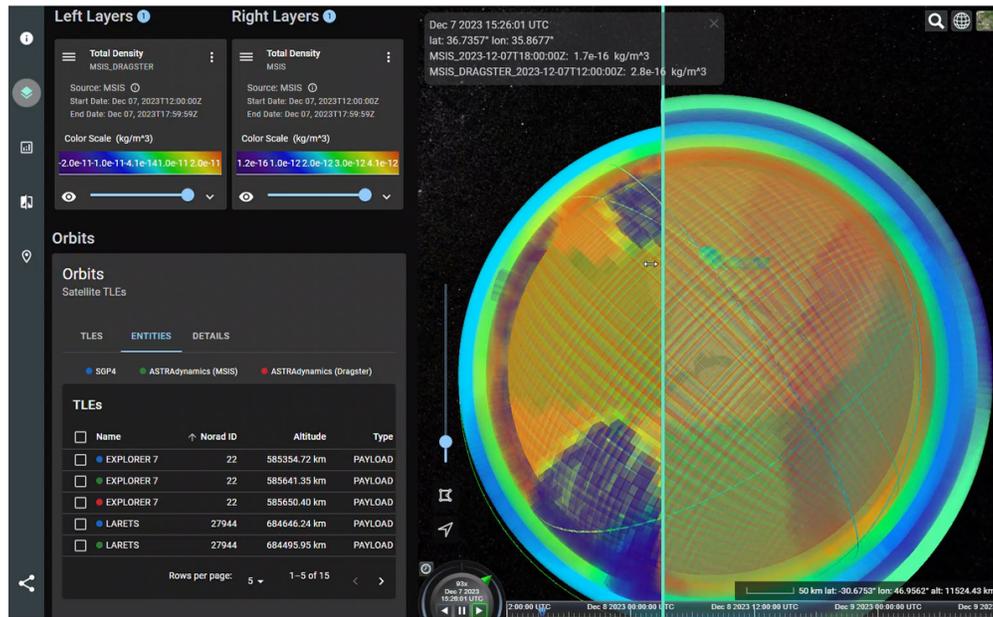


Fig. 8: Visualization of LEO satellite flying through two density models, Dragster (left) and NRLMSIS (right).

of capturing uncertainty from this better starting point. In addition, as HPC components are integrated, computational cost vs. real-time output latency will be an important trade-off to consider. Running physics-based models and large (1000+ member) ensembles of models both significantly increase the compute needed compared to empirical density models, so effectively orchestrating HPC node provisioning is an outcome of interest for commercial applications such as Dragster.

7. ACKNOWLEDGEMENTS

This work was performed under NASA SBIR Phase I and II, contract 80NSSC23CA072.

REFERENCES

- [1] Celestrak. <https://celestrak.org/SpaceData/> (accessed 2024).
- [2] CesiumJS. <https://cesium.com/platform/cesiumjs/> (accessed 2024).
- [3] GFZ. https://kp.gfz-potsdam.de/app/files/Kp_ap_Ap_SN_F107_since_1932.txt (accessed 2024).

- [4] Parallel Works Inc. <https://parallelworks.com/> (accessed 2024).
- [5] Space-Track.org. <https://www.space-track.org> (accessed 2024).
- [6] Unified Data Library. <https://unifieddatalibrary.com> (accessed 2024).
- [7] Brazilian National Institute for Space Research. *SatelliteToolbox.jl*. <https://github.com/JuliaSpace/SatelliteToolbox.jl>, May 2024.
- [8] M. Fedrizzi, T.J. Fuller-Rowell, and M.V. Codrescu. Global Joule heating index derived from thermospheric density physics-based modeling and observations. *Space Weather*, 10, 2012.
- [9] E. Kalnay. *Atmospheric Modeling, Data Assimilation and Predictability*. Cambridge University Press, 2002.
- [10] Greg Lucas. *pymis*. <https://doi.org/10.5281/zenodo.10922485>, April 2024.
- [11] J. Matzka, C. Stolle, Y. Yamazaki, O. Bronkalla, and A. Morschhauser. The geomagnetic Kp index and derived indices of geomagnetic activity. *J. Geophys. Res.*, 19, 2021.
- [12] W. Parker and R. Linares. Satellite drag analysis during the May 2024 Gannon geomagnetic storm. *AIAA Journal of Spacecraft and Rockets*, pages 1–5, 2024.
- [13] W.E. Parker, M. Freeman, G. Chisham, A. Kavanagh, P. Mun Siew, V. Rodriguez-Fernandez, and R. Linares. Influences of space weather forecasting uncertainty on satellite conjunction assessment. *Space Weather*, 22, 2024.
- [14] J.M. Picone, A.E. Hedin, D.P. Drob, and A.C. Aikin. NRLMSISE-00 empirical model of the atmosphere: Statistical comparisons and scientific issues. *J. Geophys. Res. Space Physics*, 107:SIA 15–1–SIA 15–16, 2002.
- [15] M. Pilinski, G. Crowley, M. Seaton, and E. Sutton. Dragster: An assimilative tool for satellite drag specification. *AMOS Conference Proceedings*, 2019.
- [16] M. Pilinski, G. Crowley, E. Sutton, and M. Codrescu. Improved orbit determination and forecasts with an assimilative tool for satellite drag specification. *AMOS Conference Proceedings*, 2016.
- [17] L. Qian, A.G. Burns, B.A. Emery, B. Foster, A. Lu, G. Maute, A.D. Richmond, R.G. Roble, S.C. Solomon, and W. Wang. The NCAR TIE-GCM: A community model of the coupled thermosphere/ionosphere system. *AGU Geophysical Monograph Series*, 201:73–83, 2014.