Naval Research Laboratory
Washington, DC 20375-5320

NRL/7210/MR—2022/1

Assimilative Modeling of Substorm-Initiated Traveling Ionospheric Disturbances

JOHN D. HAIDUCEK

High Angular Resolution Imaging Section
Radio/IR/Optical Sensors Branch
Remote Sensing Division

November 14, 2022

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. REPORT DATE *(DD-MM-YYYY)*	**2. REPORT TYPE**	**3. DATES COVERED** *(From - To)*
14-11-2022	NRL Memorandum Report	04-26-2021 – 04-25-2022

4. TITLE AND SUBTITLE

Assimilative Modeling of Substorm-Initiated Traveling Ionospheric Disturbances

5a. CONTRACT NUMBER

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER
NISE

6. AUTHOR(S)

John D. Haiducek

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER
N20G

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

Naval Research Laboratory
4555 Overlook Avenue, SW
Washington, DC 20375-5320

8. PERFORMING ORGANIZATION REPORT NUMBER

NRL/7210/MR--2022/1

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)

Naval Research Laboratory
4555 Overlook Avenue, SW
Washington, DC 20375-5320

10. SPONSOR / MONITOR'S ACRONYM(S)
NRL/NISE

11. SPONSOR / MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION / AVAILABILITY STATEMENT

DISTRIBUTION STATEMENT A: Approved for public release; distribution is unlimited.

13. SUPPLEMENTARY NOTES
Karles Fellowship

14. ABSTRACT

In this program, a data assimilation system was created for the NRL-developed ionosphere model SAMI3 (SAMI3 is Another Model of the Ionosphere). The aim was to enable SAMI3 to more accurately simulate traveling ionospheric disturbances (TIDs), wave-like disturbances that propagate through the Earth's ionosphere. The system assimilates total electron content (TEC) observations into an ensemble of SAMI3 simulations using an ensemble Kalman filter (enKF). The effort also led to the creation of LightDA, a general-purpose data assimilation framework which has been released as an open source project.

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			**17. LIMITATION OF ABSTRACT**	**18. NUMBER OF PAGES**	**19a. NAME OF RESPONSIBLE PERSON**
a. REPORT	**b. ABSTRACT**	**c. THIS PAGE**			John D. Haiducek
U	U	U	U	15	**19b. TELEPHONE NUMBER** *(include area code)* (202) 767-8272

Standard Form 298 (Rev. 8-98)
Prescribed by ANSI Std. Z39.18

This page intentionally left blank.

CONTENTS

FIGURES

EXECUTIVE SUMMARY

In this program, a data assimilation system was created for the NRL-developed ionosphere model SAMI3 (SAMI3 is Another Model of the Ionosphere). The aim was to enable SAMI3 to more accurately simulate traveling ionospheric disturbances (TIDs), wave-like disturbances that propagate through the Earth's ionosphere. The system assimilates total electron content (TEC) observations into an ensemble of SAMI3 simulations using an ensemble Kalman filter (enKF). The effort also led to the creation of LightDA, a general-purpose data assimilation framework which has been released as an open source project.

This page intentionally left blank

ASSIMILATIVE MODELING OF SUBSTORM-INITIATED TRAVELING IONOSPHERIC DISTURBANCES

1. BACKGROUND AND MOTIVATION

The Earth's ionosphere comprises a region of plasma located above, and partially coincident with, Earth's upper atmosphere. The ionosphere refracts and reflects radio signals that propagate through it, affecting communications, navigations, and scientific observations. Density gradients within the ionosphere can alter radio waveforms, causing degradation of signals and affecting the results of radio-based measurements. One phenomenon that contributes to such density gradients is a traveling ionospheric disturbance (TID), a wave-like structure that propagates through the ionosphere [1, 2]. TIDs can be initiated by, among other things, geomagnetic activity such as geomagnetic storms and substorms [3], which in turn are driven by solar activity including coronal mass ejections. The density gradients associated with TIDs can impact the accuracy of satellite positioning, particularly impacting applications requiring centimeter-level precision [4, 5]. In addition, TIDs affect the behavior of equatorial plasma bubbles [6, 7], which are themselves capable of disrupting satellite navigation and communications [e.g. 8].

A primary means of observing TIDs is by detecting variations in total electron content (TEC), which is the integrated electron density along a line of sight through the ionosphere. TEC observations are commonly obtained using GNSS receivers either on the ground or in space, which measure the time delays of GNSS signals induced by the ionosphere. TEC can also be measured by radio interferometers, which measure ionospherically induced phase between radio emissions from a celestial source when observed along different lines of sight. This approach has been demonstrated by NRL with the use of the Very Large Array (VLA) radio telescope to produce high-resolution images of TIDs [9]. That initial demonstration led to the creation of the VLA Low-band Ionospheric and Transient Experiment (VLITE), which provides near-continuous observations of differential TEC [10].

The aim of this project was to simulate TIDs that are induced by geomagnetic storms and substorms in SAMI3 (SAMI3 is Another Model of the Ionosphere), an NRL-developed ionospheric model [11, 12]. SAMI3 has previously been used to simulate the development of TIDs from idealized initial conditions [13] and from atmospheric drivers [14]. Developing a predictive capability for TIDs based on SAMI3 requires a more realistic model of driving parameters and initial conditions. One way to accomplish this is to use data assimilation to update a SAMI3 state to better match a set of observations. This will be based on an ensemble Kalman filter (EnKF), which updates a set of simulations to better match a set of observations [15, 16]. Data assimilation using the EnKF has been employed with other ionosphere models [17], but the application of data assimilation to SAMI3 is relatively new.

Developing capabilities to better model and predict TIDs would lead to a better understanding of their dynamics and a greater ability to plan for their effects on Navy operations that rely on satellites for navigation and communications. To do so, this research program will leverage new datasets that probe the ionosphere with higher resolution and greater precision. A new data assimilation capability will be created that will

Manuscript approved October 24, 2022.

use these datasets to update an ensemble of SAMI3 simulations to fit available observations. An additional benefit to the ensemble data assimilation approach is that the ensemble spread provides an uncertainty estimate, which would not be available from a single simulation. The end result is a simulation that is more representative of the true ionospheric state.

2. RESULTS

2.1 Data assimilation framework

To better test and validate the data assimilation algorithm, a new software library called LightDA was created that provides a modular structure for data assimilation. While several existing data assimilation frameworks are already available [18–20], the existing frameworks were found to be difficult to adapt to the problem at hand because of their size and complexity, or due to their being designed primarily to support atmospheric or ocean models rather than space environment models.

LightDA consists of a core software library that provides generic interfaces for accessing observation data and model state data, and generic interfaces to algorithms needed for the data assimilation process. Figure 1 illustrates some of the key components of LightDA. The observation set (upper right in the figure) provides a means to load observations and provides metadata required to assimilate them. The forward operator (lower right in the figure) computes predictions of observation values from a given model state. The model interface (lower left in the figure) handles reading and writing the model state, and provides additional information such as grid geometry to the forward operator and localizer. The localizer (upper right in the figure) computes localization weights, which are used to weight observations to control what parts of the model state are affected by each observation. The assimilation filter (center of the figure) computes a new model state given the prior model state, observation values, model prediction values, and localization weights.

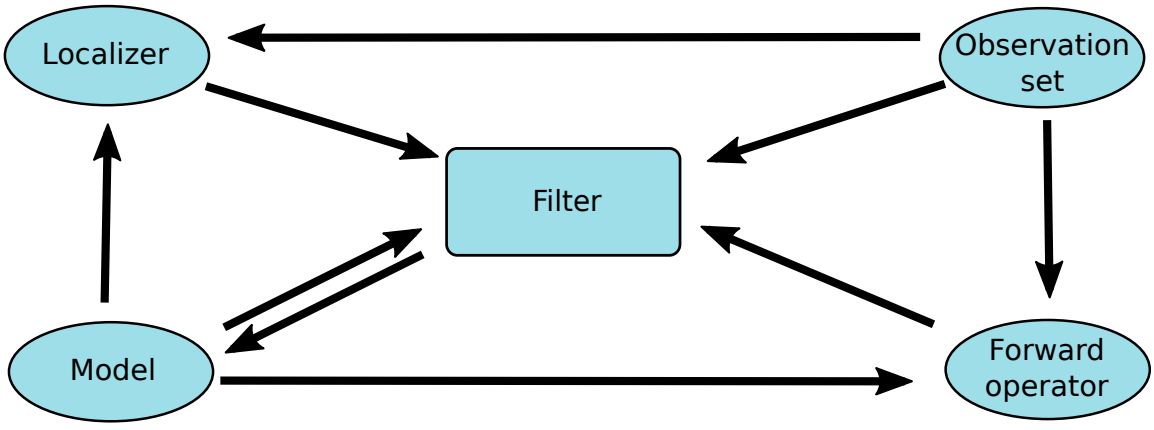

Fig. 1—Diagram showing key components of LightDA.

A key advantage of LightDA is that its modular structure enables the individual components of the data assimilation process to be developed and tested independently of one another. Tests of each component are performed using simplified minimal implementations. For example a minimal filter was created for testing

purposes that leaves the model state unchanged, providing a means to verify correct flow of data through the system. Initial functional testing was performed by interfacing LightDA to a 1-D advection solver and performing a test using synthetic observations derived from output from the same model. Figure 2 shows the results of this test. The "true" state, obtained by running the advection code with a sinusoidal initial state, is represented by a thick magenta line. Observation values, obtained by sampling the true state at selected locations and adding Gaussian noise, are shown as brown circles. The states of individual ensemble members were initialized as random noise, and in the figure they appear as colored lines overlaid on top of the true state, and the ensemble mean is shown as a black line. The ensemble mean appears close to the true state for the most part, with the largest differences between the ensemble mean and the true state appearing in places with fewer observations. The spread among ensemble members tends to be higher in these regions as well. This illustrates the usefulness of the ensemble system both for providing an estimate of the true state of the system and for providing an uncertainty estimate for the model.

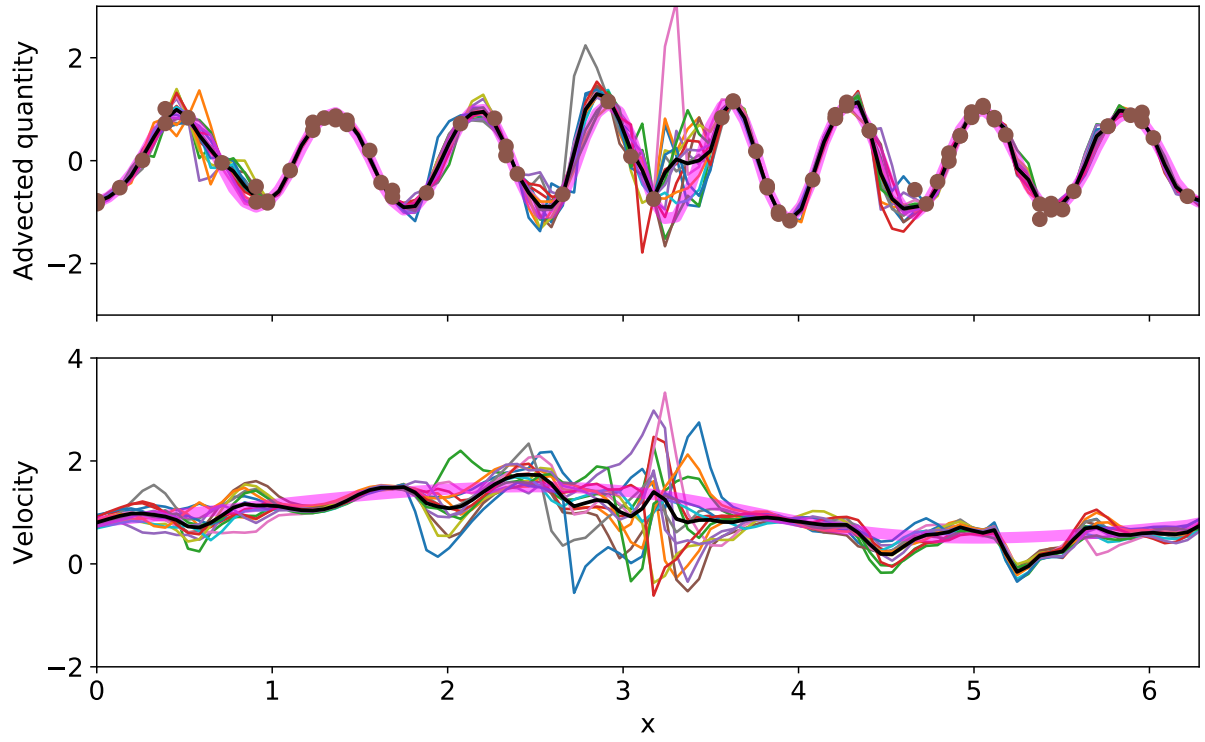

Fig. 2—Output from a test of LightDA with a 1-D advection solver. Thick magenta curve represents the "true" state of the system. States of individual ensemble members are shown as thin colored lines, with the ensemble mean depicted as a thin black line. Observation values are shown as brown circles.

2.2 Filter and localizer implementation

The filter used with both the advection solver and with SAMI3 was an EnKF implementation adapted from code found in the open source Parallel Data Assimilation Framework (PDAF) [18]. However, the parallelism in LightDA was implemented from scratch. While PDAF determines the parallelism at the model interface level, LightDA allows the model state to be loaded either in serial or in parallel, and separates the model state into chunks to be passed to the filter in parallel. For each parallel execution of the filter, a reduced

observation set is passed to the filter that includes only observations that have a nonzero localization weight for at least one point within the chunk being processed. To facilitate load balancing, the chunks are assigned to processors in random order. This random assignment reduces the chances that an uneven distribution of observations across the model domain will result in an unequal computational load among the processors. The number of elements of these chunks is configurable, and the optimal size depends on the number of processors, available memory, and characteristics of the observations.

A default localizer implementation is provided that simply returns 1 for all localization weights. For the 1-D advection solver, a localizer was implemented that computes a Gaspi-Cohn localization function [21] about each observation point. For SAMI3, a localizer was implemented that computes the Gaspi-Cohn function in magnetic longitude space about each TEC observation. No localization was performed in latitude, altitude, or L-shell in the current implementation. This was done in part for simplicity of implementation, and in part because magnetic field aligned particle transport enables coupling over long distances along field lines in the ionosphere.

2.3 Forward operator implementation

Assimilating TEC into SAMI3 requires a forward operator that computes TEC from a SAMI3 state. To do this, a procedure was written that integrates the SAMI3 electron density along a line through the SAMI3 grid. Since the data assimilation process requires computing this forward operator a fairly large number of times (equal to number of observations multiplied by the number of ensemble members), an efficient implementation is essential. Tracing through the SAMI3 grid cells requires locating each point along the line of sight to a set of three indices within the computational grid; these indices are required to obtain density data from the model state. The SAMI3 computational grid is aligned to dipole magnetic field lines, each a circle tangent to the center of the Earth. The three indices correspond to magnetic longitude, magnetic L-shell (an invariant value along the length each field line), and position along the length of the field line.

Figure 3 illustrates the SAMI3 grid, with an arrow drawn to show a notional line of sight used to compute TEC. The curved bottom line of the grid lies at a constant altitude. Field-aligned grid lines go from lower right to upper left in the figure.

Each field line has a constant magnetic longitude, and the field lines are spaced evenly along the longitude axis. This means that for a given magnetic longitude coordinate the longitude index can be computed by a simple linear map between longitude and grid indices. The other two grid indices are more complicated to find, because the grid spacing between and along magnetic field lines is uneven, and there is no analytical formula available to map grid indices to either magnetic latitude or L-shell. Thus, the only way to find the remaining two indices for a given point in physical space is to perform a search of the grid.

The initial implementation computes vertical TEC; that is, it integrates along a vertical line of sight from a point on the ground identified by latitude and longitude. This simplifies the search algorithm significantly. The longitude index only needs to be computed once since the longitude remains constant along the line of sight. The remaining integration problem is then reduced to iterating over the L-shells of the grid. This iteration begins at the L-shell that intersects the minimum altitude of the grid at the latitude of the observation, and continues through each L-shell of the grid until reaching the maximum L-shell in the computational domain. At each L-shell, a bisection search is used to find the grid index where the line of sight crosses that L-shell. At each L-shell, the electron density is interpolated to the actual intersection point, and the interpolated values are integrated using the trapezoidal rule.

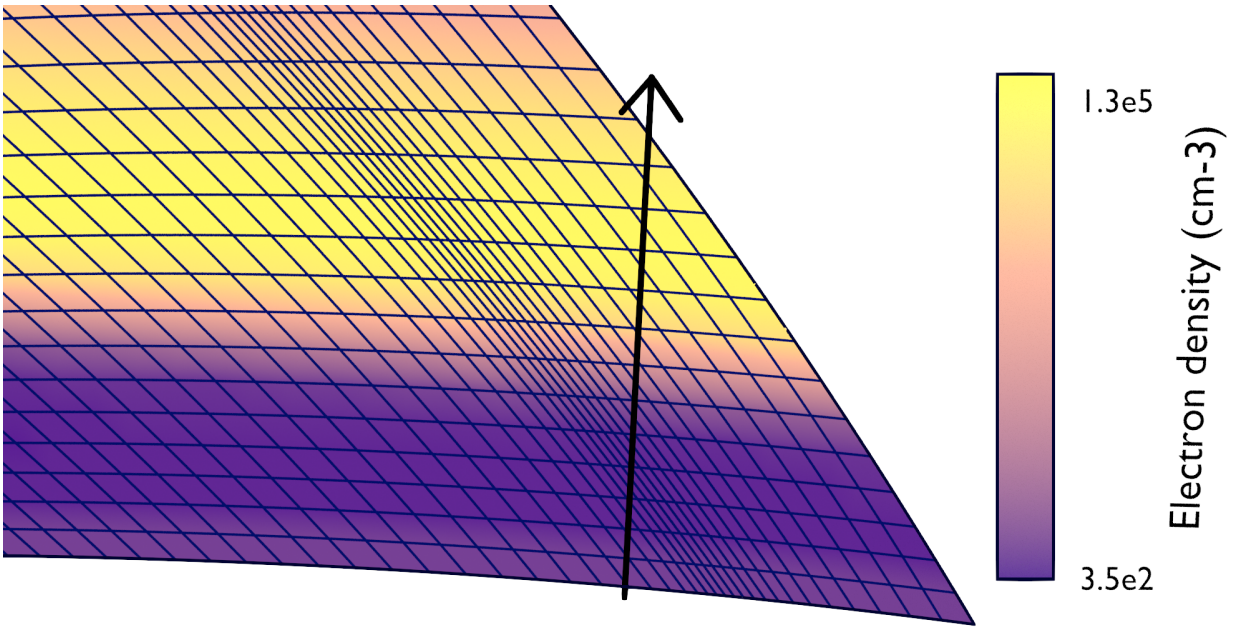

Fig. 3—Illustration of the SAMI3 grid. Colors represent electron density. The arrow illustrates a notional line of sight used to compute TEC.

Work is currently in progress to develop an algorithm for computing slant TEC. This is more complicated because the non-constant longitude along a slant path requires a three-dimensional treatment of the L-shell geometry. This could be accomplished either by computing ray-torus intersection tests of the L-shells, or by computing ray-polygon intersection tests on the cell faces from the computational grid. The ray-torus intersection problem lacks a simple analytical solution [22], but might be advantageous especially for finding the initial entry point of the ray into the grid. Ray-polygon intersection tests are comparatively straightforward, but finding the initial entry point of the grid could require multiple tests before finding the correct face. The number of tests could be reduced through the use of an acceleration tree structure.

2.4 Tests with SAMI3

The first tests of LightDA with SAMI3 aimed to verify the scalability of LightDA's parallelism scheme. This was accomplished by assimilating the same set of observations with successively larger number of processes. A standard measure of parallel scalability is speedup S, given by

$$S = \frac{t_{parallel}}{t_{serial}}, \tag{1}$$

where $t_{parallel}$ is the wall-clock execution time of the parallel algorithm, and t_{serial} is the wall-clock execution time of the same algorithm on a single processor. Ideally, speedup should increase linearly with the number of processes used, but in reality non-parallelizable portions of the code or delays in communication between processors usually result in less-than-linear speedup. In the LightDA-SAMI3 scaling test, the

problem could not be run on fewer than four processors due to memory limitations, so the execution time on four processors was used as a proxy for t_{serial}. Since t_{serial} is a constant, this affects the values of S but the shape of the curve (i.e., whether or not it is linear) is unaffected.

The results of the scaling test are shown in Figure 4. The top two panels show parallel speedup for the entire process (upper left) and for the filter alone (upper right). The lower panels show peak memory usage per node (lower left) and start-up/shut-down time (lower right).

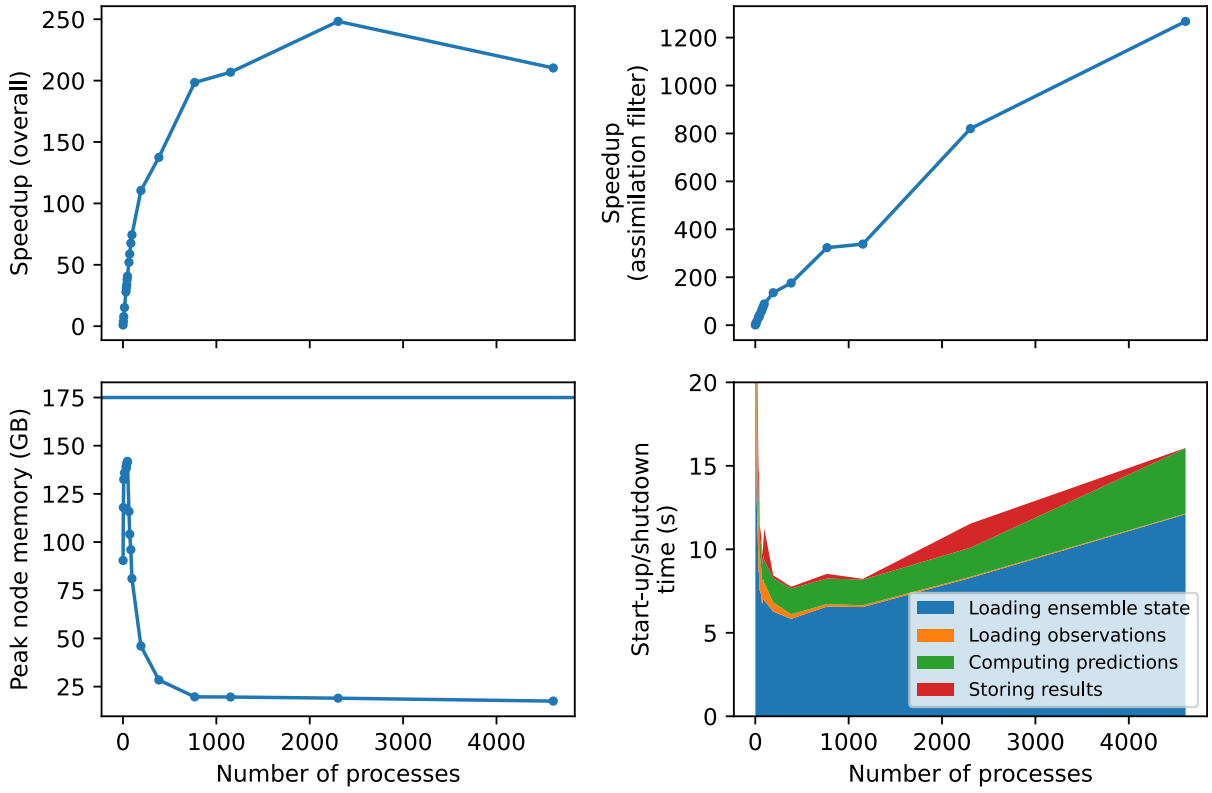

Fig. 4—Results of a scaling test to verify the LightDA parallelism scheme. Top panels: Speedup for the entire process (left) and for the EnKF filter (right). Bottom panels: Peak node memory usage (left) and time taken by start-up and shut-down tasks prior to filter execution (right).

The overall speedup (upper left panel of Figure 4) is nearly linear for the first few hundred processors, but the performance increases diminish after that, and the speedup actually decreases when the number of processes reaches 5,000. However, as shown in the upper right panel, the speedup of the EnKF filter itself remains nearly linear up to 5,000 processes. This suggests that tasks outside the filter are limiting the scalability of the system. The plot in the lower right panel confirms this. The execution time of the start-up and shut-down tasks is never much lower than 10 seconds, and even begins to increase when the number of processes is more than a few hundred. The task taking the most time (and accounting for most of the increase beyond a few hundred processors) is loading the ensemble state. This indicates that scaling the system beyond a few hundred processes will require increasing the parallelism of the model i/o interface so that the model state can be loaded in a more parallel manner. However, the existing parallel performance has been found to be acceptable for current needs.

The upper panel of Figure 5 shows root mean squared error (RMSE) for TEC. The blue curve shows RMSE for the SAMI3 model output, and the orange curve shows the TEC values obtained from the ensemble after running the EnKF filter. The filter consistently maintains an RMSE of less than 5. The lower panel of Figure 5 shows the ensemble standard deviation of the TEC observations before and after executing the EnKF filter. The standard deviation is quite low, generally between 0.03 and 0.04. With an EnKF, the ensemble spread should normally be comparable to or larger than the observation errors. A smaller ensemble spread results in the filter making only small changes to the model state. Efforts are ongoing to find appropriate settings that will maintain a wider ensemble spread to enable more effective assimilation.

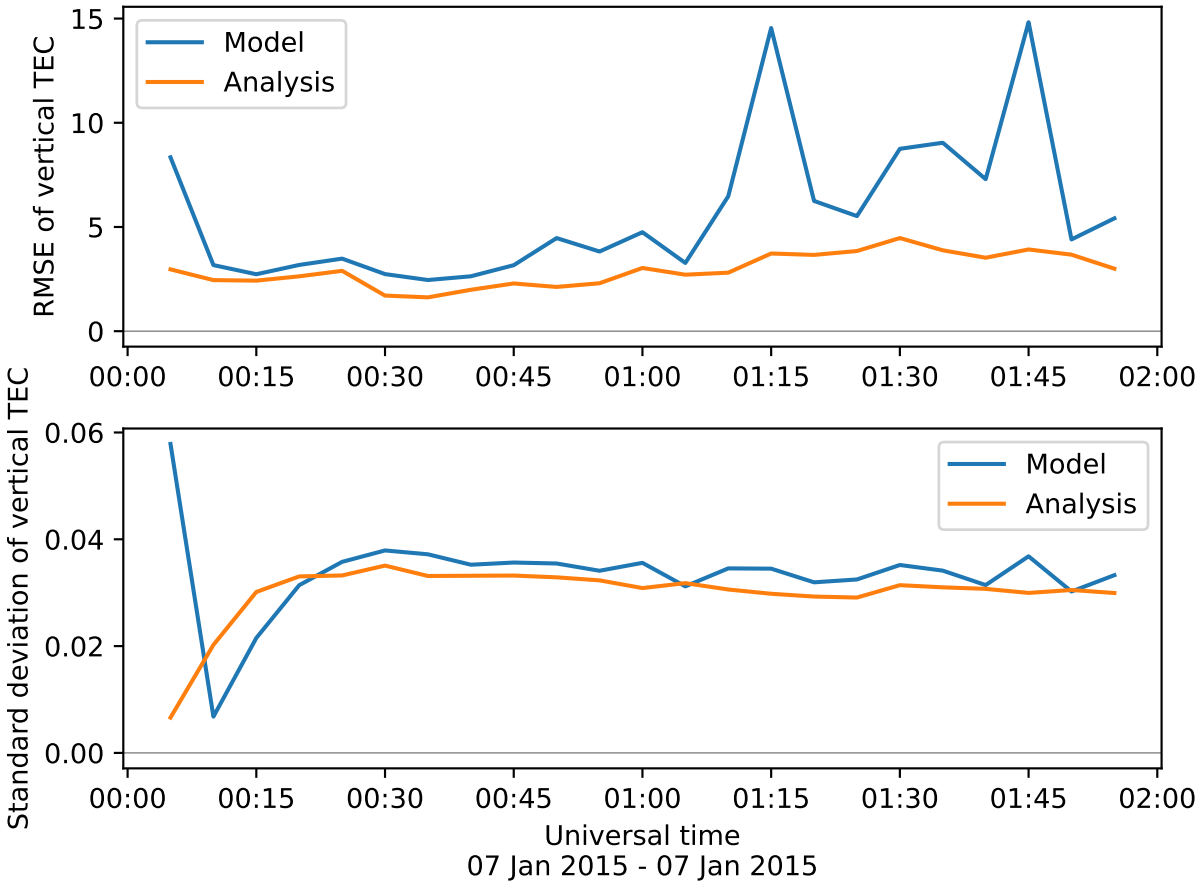

Fig. 5—Top: Root mean squared error (RMSE) of LightDA-SAMI3 relative to TEC observations. Bottom: Ensemble standard deviation of TEC predictions.

3. NEXT STEPS

With LightDA's basic capabilities having now been tested and verified, the immediate next step will be to publish a journal article describing LightDA. Another immediate next step is to continue refining the assimilation of TEC into SAMI3. Currently the initial ensemble is perturbed in electron density alone, and this has been identified as a possible contributing factor to the narrow ensemble spread produced by the filter. Another factor is that all ensemble members currently use the same value of F10.7, a solar radio flux parameter used as a proxy for ultraviolet flux. Solar ultraviolet radiation is a main driver of ionization, and

as a result F10.7 is a major influence on ionization levels in SAMI3. Varying F10.7 among the ensemble members, and possibly updating F10.7 with the filter, may lead to a larger ensemble spread.

The forward operator that computes TEC is currently only capable of computing vertical TEC. This was done for simplicity, but most real-world TEC observations are taken at an angle from the vertical, and vertical TEC is estimated from the original slant TEC observations. Assimilating slant TEC directly will provide information about the vertical distribution of plasma in the ionosphere. Developing the procedure to trace through the grid at angles also opens up the possibility of assimilating radio occultation (RO) observations obtained by satellites. RO observations are taken on lines of sight from positions in Earth orbit toward GNSS satellites, and estimating them with SAMI3 requires tracing through the grid at an angle similarly to what is required to compute slant TEC from the ground. In either case, the implementation will require implementing a search in longitude and L-shell simultaneously. While the linear mapping from magnetic longitude to grid index still applies, the longitude is not known a priori as in the vertical TEC case. Instead, an intersection point must be found for each L-shell, either using a ray-torus intersection algorithm or by performing ray-polygon intersection tests on the cell faces from the grid.

Ionosondes provide yet another potential dataset to be assimilated. An ionosonde is a specialized radar tranceiver used to measure plasma density at various altitudes directly above the instrument. In this case, the vertical tracing algorithm can be used, but it would need to be modified to output electron density values along the line of sight rather than the integrated TEC value.

In the process of developing LightDA and creating interfaces for SAMI3, a number of areas for architectural improvement to LightDA have been identified. The current implementation of LightDA requires the various components to be chosen and configured at compile time. A more flexible design would allow these to be defined in a configuration file that would be read when the application starts up. This would make the system more flexible and easier to use. At the same time, the parallelism of LightDA can be improved to be more flexible. The current implementation avoids making assumptions about how the model is parallelized, but it does impose assumptions about the parallelism of other components. For example, the filter is assumed to be serial so that the assimilation manager can manage the parallelism of the filter. A more flexible design would allow for the filter to implement its own parallelism scheme, with the current chunking scheme provided as an option for serial filter implementations. Finally, a generic interface can be provided to allow for additional operations to be added to the system; for example, to add components that modify the ensemble state for the purpose of adding perturbations to the ensemble, increasing the spread between ensemble members, or for fixing problems with the filter output (e.g. replacing negative density values with positive ones). These kinds of operations are currently added to the code in somewhat arbitrary ways, and having a standard way of handling them would provide much greater flexibility and make these operations more transparent to the user.

The software developed under this Karles fellowship has been mentioned newly funded ONR proposal, "Thermospheric/Ionospheric Data Assimilation." This program aims to develop an operational data assimilation system for the ionosphere and thermosphere. Under this program, the data assimilation tools developed here will be combined with a local ensemble transform Kalman filter (LETKF) developed by Code 7227, to support data assimilation both with SAMI3 and with thermospheric models for operational use by the Navy. It is expected that the forward operator developed in this effort will be used in this new program. It is not yet clear whether the LightDA core library will be used directly as part of this effort, but the ideas and experience gained in the process of its development will contribute to the efforts to develop an operational data assimilation system for Navy use, and LightDA's simplicity and flexibility make it an ideal platform for future data assimilation experiments.

REFERENCES

1. M. J. Davis and A. V. da Rosa, "Traveling ionospheric disturbances originating in the auroral oval during polar substorms," *Journal of Geophysical Research* **74**(24), 5721–5735 (nov 1969), ISSN 01480227, doi:10.1029/JA074i024p05721. URL http://doi.wiley.com/10.1029/JA074i024p05721.

2. R. Leitinger, R. Leitinger, and M. Rieger, "The TID model for modulation of large scale electron density models," *Annals of Geophysics* **48**(3) (dec 2005), ISSN 1593-5213, doi:10.4401/ag-3216. URL http://www.annalsofgeophysics.eu/index.php/annals/article/view/3216.

3. G. Kirchengast, "Characteristics of high-latitude TIDs from different causative mechanisms deduced by theoretical modeling," *Journal of Geophysical Research: Space Physics* **102**(A3), 4597–4612 (mar 1997), ISSN 01480227, doi:10.1029/96JA03294. URL http://doi.wiley.com/10.1029/96JA03294.

4. M. Hernández-Pajares, J. M. Juan, J. Sanz, and O. L. Colombo, "Application of ionospheric tomography to real-time GPS carrier-phase ambiguities resolution, at scales of 400-1000 km and with high geomagnetic activity," *Geophysical Research Letters* **27**(13), 2009–2012 (jul 2000), ISSN 00948276, doi:10.1029/1999GL011239. URL http://doi.wiley.com/10.1029/1999GL011239.

5. M. Hernández-Pajares, J. M. Juan, and J. Sanz, "Medium-scale traveling ionospheric disturbances affecting GPS measurements: Spatial and temporal analysis," *Journal of Geophysical Research* **111**(A7), A07S11 (jul 2006), ISSN 0148-0227, doi:10.1029/2005JA011474. URL http://doi.wiley.com/10.1029/2005JA011474.

6. S. Singh, F. S. Johnson, and R. A. Power, "Gravity wave seeding of equatorial plasma bubbles," *Journal of Geophysical Research: Space Physics* **102**(A4), 7399–7410 (apr 1997), ISSN 01480227, doi:10.1029/96JA03998. URL http://doi.wiley.com/10.1029/96JA03998.

7. Y. Otsuka, K. Shiokawa, and T. Ogawa, "Disappearance of equatorial plasma bubble after interaction with mid-latitude medium-scale traveling ionospheric disturbance," *Geophysical Research Letters* **39**(14) (jul 2012), ISSN 00948276, doi:10.1029/2012GL052286. URL http://doi.wiley.com/10.1029/2012GL052286.

8. S. Basu, E. MacKenzie, and S. Basu, "Ionospheric constraints on VHF/UHF communications links during solar maximum and minimum periods," *Radio Science* **23**(3), 363–378 (may 1988), ISSN 00486604, doi:10.1029/RS023i003p00363. URL http://doi.wiley.com/10.1029/RS023i003p00363.

9. J. F. Helmboldt, "Drift-scan imaging of traveling ionospheric disturbances with the Very Large Array," *Geophysical Research Letters* **41**(14), 4835–4843 (jul 2014), ISSN 00948276, doi:10.1002/2014GL060951. URL http://doi.wiley.com/10.1002/2014GL060951.

10. T. E. Clarke, N. E. Kassim, J. F. Helmboldt, P. S. Ray, W. M. Peters, B. Hicks, W. Brisken, R. A. Perley, F. N. Owen, and H. Intema, "Low Frequencies on the NRAO VLA and the new VLA Ionospheric and Transient Experiment (VLITE)," *American Astronomical Society, AAS Meeting #225, id.311.04* **225** (2015). URL http://adsabs.harvard.edu/abs/2015AAS...22531104C.

11. J. D. Huba, G. Joyce, and J. A. Fedder, "Sami2 is Another Model of the Ionosphere (SAMI2): A new low-latitude ionosphere model," *Journal of Geophysical Research: Space Physics* **105**(A10), 23035–23053 (oct 2000), ISSN 01480227, doi:10.1029/2000JA000035. URL http://doi.wiley.com/10.1029/2000JA000035.

12. J. D. Huba, G. Joyce, and J. Krall, "Three-dimensional equatorial spread *F* modeling," *Geophysical Research Letters* **35**(10) (may 2008), ISSN 00948276, doi:10.1029/2008GL033509. URL http://doi.wiley.com/10.1029/2008GL033509.

13. T. M. Duly, J. D. Huba, and J. J. Makela, "Self-consistent generation of MSTIDs within the SAMI3 numerical model," *Journal of Geophysical Research: Space Physics* **119**(8), 6745–6757 (aug 2014), ISSN 21699380, doi:10.1002/2014JA020146. URL http://doi.wiley.com/10.1002/2014JA020146.

14. J. D. Huba, D. P. Drob, T. W. Wu, and J. J. Makela, "Modeling the ionospheric impact of tsunami-driven gravity waves with SAMI3: Conjugate effects," *Geophysical Research Letters* **42**(14), 5719–5726 (jul 2015), ISSN 00948276, doi:10.1002/2015GL064871. URL http://doi.wiley.com/10.1002/2015GL064871.

15. G. Evensen, "Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics," *Journal of Geophysical Research* **99**(C5), 10143 (may 1994), ISSN 0148-0227, doi:10.1029/94JC00572. URL http://doi.wiley.com/10.1029/94JC00572.

16. G. Evensen, "The Ensemble Kalman Filter: theoretical formulation and practical implementation," *Ocean Dynamics* **53**(4), 343–367 (nov 2003), ISSN 1616-7341, doi:10.1007/s10236-003-0036-9. URL http://link.springer.com/10.1007/s10236-003-0036-9.

17. G. A. Hajj, B. D. Wilson, C. Wang, X. Pi, and I. G. Rosen, "Data assimilation of ground GPS total electron content into a physics-based ionospheric model by use of the Kalman filter," *Radio Science* **39**(1) (feb 2004), ISSN 00486604, doi:10.1029/2002RS002859. URL https://trs.jpl.nasa.gov/bitstream/handle/2014/7673/03-1803.pdf?sequence=1http://doi.wiley.com/10.1029/2002RS002859.

18. L. Nerger, W. Hiller, and J. Schröter, "PDAF - the Parallel Data Assimilation Framework: Experiences with Kalman Filtering," in W. Zwieflhofer and G. Mozdzynski, eds., *Use of High Performance Computing in Meteorology - Proceedings of the 11. ECMWF Workshop*, pp. 63–83 (WORLD SCIENTIFIC, Sept. 2005), ISBN 978-981-256-354-5, doi:10.1142/9789812701831_0006.

19. J. Anderson, T. Hoar, K. Raeder, H. Liu, N. Collins, R. Torn, and A. Avellano, "The Data Assimilation Research Testbed: A Community Facility," *Bulletin of the American Meteorological Society* **90**(9), 1283–1296 (Sept. 2009), ISSN 0003-0007, doi:10.1175/2009BAMS2618.1.

20. Y. Trémolet and T. Auligné, "The Joint Effort for Data Assimilation Integration (JEDI)," *JCSDA Quarterly Newsletter* pp. 1–5 (2020), doi:10.25923/RB19-0Q26.

21. G. Gaspari and S. E. Cohn, "Construction of Correlation Functions in Two and Three Dimensions," *Quarterly Journal of the Royal Meteorological Society* **125**(554), 723–757 (1999), ISSN 1477-870X, doi:10.1002/qj.49712555417.

22. V. Skala, "Line-Torus Intersection for Ray Tracing: Alternative Formulations **12**(7), 10 (2013).