Data assimilation and driver estimation for the Global Ionosphere–Thermosphere Model using the Ensemble Adjustment Kalman Filter

Alexey V. Morozov [a,*], Aaron J. Ridley [a], Dennis S. Bernstein [a], Nancy Collins [b], Timothy J. Hoar [b], Jeffrey L. Anderson [b]

[a] University of Michigan, Ann Arbor, MI, United States
[b] National Center for Atmospheric Research, Boulder, CO, United States

ABSTRACT

This paper proposes a differential inflation scheme and applies this technique to driver estimation for the Global Ionosphere–Thermosphere Model (GITM) using the Ensemble Adjustment Kalman Filter (EAKF), which is a part of the Data Assimilation Research Testbed (DART). Driver estimation using EAKF is first demonstrated on a linear example and then applied to GITM. The Challenging Minisatellite Payload (CHAMP) neutral mass density measurements are assimilated into a biased version of GITM, and the solar flux index, $F_{10.7}$, is estimated. Although the estimate of $F_{10.7}$ obtained using DART does not converge to the measured values, the converged values are shown to drive the GITM output close to CHAMP measurements. In order to prevent the ensemble spread from converging to zero, the state and driver estimates are inflated. In particular, the $F_{10.7}$ estimate is inflated to have a constant variance. It is shown that EAKF with differential inflation reduces the model bias from 73% down to 7% along the CHAMP satellite path when compared to the biased GITM output obtained without using data assimilation. The Gravity Recovery and Climate Experiment (GRACE) density measurements are used to validate the data assimilation performance at locations different from measurement locations. It is shown that the bias at GRACE locations is decreased from 76% down to 52% as compared to not using data assimilation, showing that model estimation of the thermosphere is improved globally.

© 2013 Elsevier Ltd. All rights reserved.

1. Introduction

Solar radiation is the largest source of energy in the thermosphere, significantly affecting the neutral mass density, ρ, which in turn influences the drag experienced by objects in low-Earth orbit. Uncertainty in drag translates to uncertainty in position that may even result in a loss of spacecraft (Storz et al., 2005). One way of decreasing drag uncertainty is by obtaining more precise estimates of the neutral density from thermospheric models. An effective way of improving the accuracy of these models is data assimilation.

For example, assimilation of azimuth, elevation, range and range-rate measurements from calibration satellites into the high-accuracy satellite drag model (HASDM) resulted in about 25% bias reduction as compared to the case with no data assimilation (Storz et al., 2005).

Using an alternative data source, Matsuo et al. (2012) showed that assimilating Challenging Minisatellite Payload (CHAMP,

Reigber et al., 2002, orbit altitude of about 400 km) and Gravity Recovery and Climate Experiment (GRACE, Tapley et al., 2004, orbit altitude of about 500 km) mass density measurements resulted in about 50% bias reduction for the Coupled Thermosphere Ionosphere Plasmasphere Electrodynamics (CTIPe, Millward et al., 2001) model, as compared to the no data assimilation case. A modified Kalman filter was used to perform data assimilation with a prior error covariance matrix constructed using empirical orthogonal functions with a maximum-likelihood update method.

CTIP (Millward et al., 1996) was also used by Kersley et al. (2005), which interpreted the results of radio tomographic imaging performed on the U.S. Navy Ionospheric Measuring System and compared these results with GPS total electron content (TEC, Mannucci et al., 1998) measurements. The authors concluded that, while the small-scale details within the data were beyond the capabilities of the model, the general features captured by the model aided the interpretation of the tomography results.

Similarly, Utah State University's Global Assimilation of Ionospheric Measurements (USU GAIM, Schunk et al., 2004; Scherliess et al., 2004) framework also demonstrated how initial model bias

---

* Corresponding author.
E-mail address: morozova@umich.edu (A.V. Morozov).
URL: http://umich.edu/~morozova (A.V. Morozov).

can be removed using an approximate Kalman filter. The Kalman filter was used to assimilate the slant TEC measurements and ionosonde density profiles into the ionosphere plasmasphere model (IPM) to estimate one of the model drivers, namely, the equatorial vertical plasma drift. A reduced state Kalman filter was implemented through numerical linearization of reduced-state IPM at each time step. The authors concluded that, while this linearization might require thousands of model runs, each model run is parallel and computationally efficient.

Likewise, the performance of a low-resolution version of the Jet Propulsion Laboratory/University of Southern California Global Assimilation Ionospheric Model (JPL/USC GAIM, Pi et al., 2003; Wang et al., 2004) is described by Mandrake et al. (2005). The GPS slant TEC measurements from 200 ground GPS sites were used to estimate the single-ion 3D density state using a band-limited Kalman filter approximation. The results of the assimilation were validated against the withheld ground measurements, and the vertical TEC measurements from the JASON satellite (Perbos et al., 2003). The study concluded that, when apparent bias in the JASON measurements was removed, the GAIM performance was improved from 7 total electron content units (TECU) to about 5.3 TECU.

An application of data assimilation for the Global Ionosphere–Thermosphere Model (GITM) was described by Kim et al. (2008); Kim (2008). In particular, Kim et al. (2008) used the localized unscented Kalman filter (LUKF) to assimilate electron number density, ion temperature, and ion velocity into the one-dimensional (1D) GITM, described in detail in Ridley et al. (2006). That study concluded that using a small localization region (about 11 altitude cells out of 50) in LUKF is comparable to the nominal unscented Kalman filter (UKF) performance without the added computational cost. Kim (2008) also demonstrated that the GITM data assimilation could be extended to include three-dimensional (3D) GITM using the ensemble Kalman filter (EnKF, Evensen, 1994). More precisely, 7 ensemble members were used in the EnKF to assimilate the simulated electron number density, ion temperature, neutral density, and (separately) TEC at six fixed locations.

The present study, reported here, extended the preliminary findings of Kim (2008) by increasing the ensemble size and adding driver estimation. Driver estimation is sought to improve bias removal in the GITM neutral mass density estimates as compared to the CHAMP measurements. Following the example of Matsuo et al. (2012), the GRACE measurements were used to validate the results, though, in this case, the GRACE neutral density measurements were not scaled and were used at the altitude of measurements (approximately 498 km). The CHAMP measurements were assimilated into GITM via the ensemble adjustment Kalman filter (EAKF, Anderson, 2001). EAKF was chosen over the EnKF as it does not perturb measurements but otherwise provides similar performance. The EAKF, implemented in the Data Assimilation Research Testbed (DART), which was written and is maintained by the National Center for Atmospheric Research (Anderson et al., 2009), was used in this study. DART includes the EnKF and EAKF, as well as several other filters and has been used in numerical weather prediction (Li and Liu, 2009; Stensrud and Gao, 2010), aerosol modeling systems (Khade et al., 2012), and upper atmosphere modeling (Lee et al., 2012).

This paper is organized as follows. Section 2.1 introduces GITM and demonstrates the need for data assimilation by highlighting the mismatch between ρ produced by GITM and ρ measured by CHAMP when bias is added to GITM. EAKF is introduced independently of GITM in Sections 2.2–2.4 to separate general data assimilation concepts from model-dependent concepts. Section 2.5 addresses localization. The results of performing data assimilation with GITM are shown in Section 3.

2. Methodology

2.1. Global ionosphere–thermosphere model

GITM is a global model of the upper atmosphere, whose state variables include neutral and ion densities, temperatures, and velocities, as well as electron density and temperature (Ridley et al., 2006). One notable feature of GITM is that it produces nonhydrostatic solutions (Deng et al., 2008) by solving the vertical momentum equation for each neutral species. The species momenta are then coupled using friction terms such that in the lower thermosphere the major species (for example, N_2) tend to force the minor species (for example, NO, O) to be nonhydrostatic. The relaxation of the hydrostatic assumption is useful for modeling the auroral region due to the presence of strong heating effects. The electric potential and the average energy used in polar region of GITM come from the electric potential model described in Weimer (1996) and the auroral model described in Fuller-Rowell and Evans (1987), respectively. Additionally, GITM's lower boundary conditions are specified by the model based on the mass spectrometer and incoherent scatter data (MSIS, Hedin, 1987) and the international reference ionosphere (IRI, Bilitza, 2001; for more details on the coupling, see Ridley et al., 2006).

One of the main differences between the upper atmosphere and troposphere is the effect of drivers. The time evolution of the troposphere is determined mainly by its initial conditions. In the upper atmosphere, however, the highly time-dependent drivers play a more important role, thereby making the upper atmosphere a contractive system (Sontag, 2010; Russo et al., 2010). For example, the relationship between neutral temperature and solar radio flux at the wavelength of 10.7 cm ($F_{10.7}$) can be shown to have a positive correlation. This implies that, as $F_{10.7}$ increases, the neutral temperature will increase in the atmosphere region closest to the sun. Various studies (Pawlowski and Ridley, 2008) have shown that, when the solar flux increases dramatically (as in a solar flare), the density on the day side responds rapidly; in contrast, the rest of the atmosphere responds with a time delay as the perturbation wave propagates to the night side. It can also be shown that, as temperature increases, the volume of the atmosphere increases and pushes the upper atmosphere higher in altitude. Accordingly, a satellite orbiting at a fixed altitude would observe an increase in ρ on the day side. The direct relationship between ρ and $F_{10.7}$ is exploited in Section 3.1, where the density at about 400 km above the subsolar point is assimilated to estimate $F_{10.7}$.

While models can provide insight into the dynamics of the upper atmosphere, it is often necessary to validate and calibrate the model dynamics using measurements. The CHAMP satellite, which orbited from 2000 to 2010, provided measurements of the upper atmospheric state. One measurement was of the acceleration experienced by the satellite, which can be used to infer drag. Since drag is proportional to the mass density of the thermosphere, mass density can be estimated from CHAMP measurements (Sutton, 2009; Sutton et al., 2007). Neutral density readings from CHAMP are separated in time by approximately 47 s on an average. The CHAMP trajectory (87.3° inclination with an average altitude of about 400 km during the simulated days) is shown in Fig. 1(a). Fig. 1(b) shows that version of GITM with no data assimilation ($F_{10.7} = 142$ SFU) underestimates ρ by about 2×10^{-12} kg m^{-3} compared to CHAMP measurements between 01 UT and 04 UT. One reason for this mismatch is that the high-latitude drivers (electric potential and aurora) are intentionally held steady during this time period to cause GITM temperatures to drop, thereby forcing a bias in GITM neutral density. It is shown in Yiğit and Ridley (2011) that variations in high-latitude drivers can cause heating. Therefore, by purposefully holding the drivers
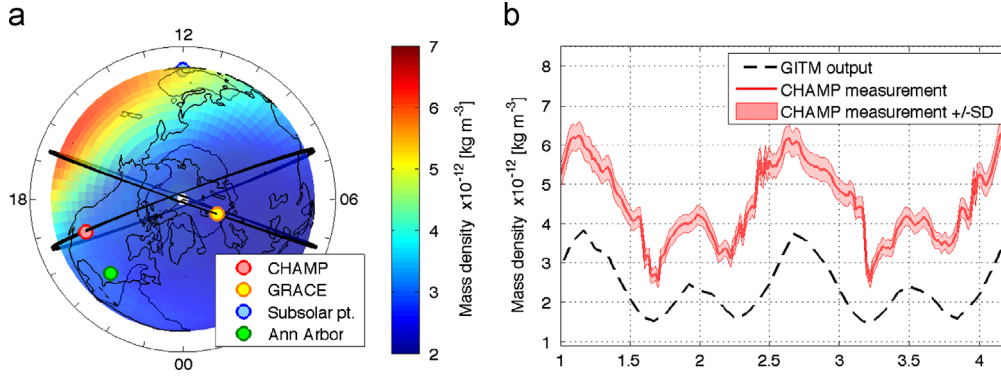
Fig. 1. (a) shows the neutral density at about 400 km altitude at 02:32 UT 01/12/2002 as well as CHAMP and GRACE orbital trajectories. CHAMP (red) and GRACE (yellow) orbits for the chosen time period do not pass near the subsolar point. The GITM gridpoint closest to Ann Arbor, MI (green) is referred to in subsequent figures. (b) demonstrates the model-data mismatch between no data assimilation GITM output (black) and CHAMP measurement (red). GITM without data assimilation underestimates neutral density compared to CHAMP measurements. (For interpretation of the references to color in this figure caption, the reader is referred to the web version of this paper.)

steady, the atmosphere cools off, and the neutral density at CHAMP altitudes decreases.

The following sections establish how model bias can be removed by estimating drivers, whereas Section 3 applies this technique and demonstrates how GITM estimate of neutral density at CHAMP locations can be brought closer to CHAMP measurements by estimating $F_{10.7}$.

2.2. EAKF: data assimilation without driver estimation

This section demonstrates data assimilation in a simple and illustrative context, namely as applied to a first order, scalar linear system. The system parameters are chosen randomly and are not intended to represent GITM. The goal of the following experiment is to estimate the state (s_k) based on the noisy measurements (y_k) without knowledge of the driver (u_k). Consider the linear system

$$s_k = 0.5 s_{k-1} + u_{k-1}, \quad u_k = 1.0 + \sin(0.4k), \tag{1}$$

$$y_k = s_k + v_k, \quad v_k \sim N(0, \ 0.3), \tag{2}$$

where k is the time step, s_k is the true state of the system, u_k is the unknown time-varying driver, y_k is the state measurement corrupted by noise v_k sampled from a zero-mean normal distribution with known variance $R_k = 0.3$. For the ith ensemble member, $\hat{s}_{k,i}^+$ denotes the posterior state estimate just after the measurement y_k is assimilated, whereas $\hat{s}_{k,i}^-$ denotes the prior state estimate just before assimilating y_k. The prior state estimate is calculated from the posterior estimate at the previous step by propagating the dynamics, as given by

$$\hat{s}_{k,i}^- = 0.5 \hat{s}_{k-1,i}^+ + \hat{u}_{k-1,i}, \tag{3}$$

where, for each ensemble member, $\hat{u}_{k-1,i}$ is a constant value selected randomly from a normal distribution (driver estimation is explored in Section 2.3). EAKF assumes that $\hat{s}_{k,i}^-$ and y_k are normal random variables whose statistics can be described by their means and variances. The prior state variance can thus be approximated by the sample variance

$$\sigma^2[\hat{s}_k^-] = \sum_{i=1}^N (\hat{s}_{k,i}^- - \mu[\hat{s}_k^-])^2/(N-1), \tag{4}$$

where $\mu[\hat{s}_k^-]$ is the prior ensemble sample mean. EAKF generates the posterior ensemble estimate by applying Bayes theorem. Accordingly, the posterior probability density function (PDF) is equal to the normalized product of the prior and the measurement likelihood (Anderson, 2001). Since the product of two normal random variables is also a normal random variable, the posterior state estimate PDF is also normal, and its mean and variance are

given by

$$\mu[\hat{s}_k^+] = \sigma^2[\hat{s}_k^+] \left[\frac{\mu[\hat{s}_k^-]}{\sigma^2[\hat{s}_k^-]} + \frac{y_k}{R_k} \right] \tag{5}$$

$$\mu[\hat{s}_k^+] = \frac{R_k}{\sigma^2[\hat{s}_k^-] + R_k} \mu[\hat{s}_k^-] + \frac{\sigma^2[\hat{s}_k^-]}{\sigma^2[\hat{s}_k^-] + R_k} y_k, \tag{6}$$

$$\sigma^2[\hat{s}_k^+] = \left[\frac{1}{\sigma^2[\hat{s}_k^-]} + \frac{1}{R_k} \right]^{-1}. \tag{7}$$

The final step of EAKF scales and shifts the prior ensemble to match the new mean and variance, as given by

$$\hat{s}_{k,i}^+ = \sqrt{\frac{\sigma^2[\hat{s}_k^+]}{\sigma^2[\hat{s}_k^-]}} \left(\hat{s}_{k,i}^- - \mu[\hat{s}_k^-] \right) + \mu[\hat{s}_k^+]. \tag{8}$$

To demonstrate the overall performance of (1)–(8), 20 ensemble members are initialized from a zero-mean normal initial distribution with variance of 0.4 for both the state and driver estimates. The initial variance is chosen to be greater than the noise variance (R_k) so as to initially weight the measurements y_k more heavily. More precisely, $\sigma^2 > R_k$ implies that the measurement weight $\sigma^2/(\sigma^2 + R_k)$ is greater than the prior mean weight $R_k/(\sigma^2 + R_k)$ in (6). Eqs. (1)–(8) are propagated forward in time for 50 steps, while the measurements y_k are assimilated at every step.

Fig. 2 shows the resulting state and driver estimates as a function of time. Along with plotting the true state (s_k), Fig. 2 (a) shows the measurement (y_k), the initial ensemble distribution for the posterior state estimate ($\hat{s}^+$), and the mean and spread of the ensemble ($\mu[\hat{s}]$ and $\sigma[\hat{s}]$). The line depicting the ensemble mean (blue) consists of prior and posterior estimate, thereby resulting in a discernible discontinuity at every step. Fig. 2(a) demonstrates that, although the posterior mean approaches the true state s_k, the prior mean deviates away from s_k due to an incorrect driver value used by the ensemble members.

The performance can be better quantified in terms of estimation errors. The prior state root mean square percentage error (RMSPE) is defined as

$$\text{RMSPE}^- = 100 \times \frac{\sqrt{\mu[(\mu[\hat{s}_k^-] - s_k)^2]}}{\sqrt{\mu[s_k^2]}}, \tag{9}$$

and the posterior RMSPE is defined by replacing the minus superscripts by plus superscripts. Calculating RMSPE for the estimates shown in Fig. 2(a) results in prior and posterior state RMSPE values of 50.92% and 26.41%, respectively. To decrease the

Fig. 2. Time evolution of the EAKF state estimate for the system (1), (2) without driver estimation. (a) shows the true state (black dashed), the state measurement (red), the state estimate initial distribution (green), and the state estimate mean (dark blue) and spread (light blue, standard deviation). (b) shows the true driver, driver estimate distribution, mean, and variance in the same colors as (a). The posterior ensemble mean converges to the true state, whereas the prior mean deviates during the model propagation step due to the incorrect driver estimate. (For interpretation of the references to color in this figure caption, the reader is referred to the web version of this paper.)

influence of a particular noise realization, this experiment was repeated using different noise realizations 1000 times, and the averaged prior state RMSPE was found to be 62.40%. The two-fold decrease between state prior and posterior errors means that either the model is inherently incorrect or that the drivers are incorrect, since the model keeps forcing the results away from the data. This implies that such a model would have a poor forecasting ability. The next section shows how the inconsistency between the prior and posterior states can be reduced by estimating the driver.

2.3. EAKF: data assimilation with driver estimation

To perform driver estimation, the augmented state vector is defined as $x_k = [s_k; u_k]$. Since state augmentation makes the problem of estimating x_k multivariate, the EAKF equations take on a matrix form, as derived in Anderson (2001).

The matrix formulation of EAKF can be described in joint state-measurement space (Anderson, 2001; Tarantola, 2005), where the joint state vector is defined as $z_k = [x_k; y_k] = [s_k; u_k; y_k]$ Then, (3)–(8) for the ith ensemble member become

$$\hat{z}^-_{k,i} = [0.5\hat{s}^+_{k-1,i} + \hat{u}^+_{k-1,i}; \hat{u}^+_{k-1,i}; \hat{s}^+_{k,i}], \tag{10}$$

$$P^-_k = \sum_{i=1}^{N} (\hat{z}^-_{k,i} - \mu[\hat{z}^-_k])(\hat{z}^-_{k,i} - \mu[\hat{z}^-_k])^T / (N-1), \tag{11}$$

$$\mathcal{A}_k = (\mathcal{F}^T_k)^{-1}\mathcal{G}^T_k(U^T_k)^{-1}B^T_k(\mathcal{G}^T_k)^{-1}\mathcal{F}^T_k, \tag{12}$$

$$P^+_k = [(P^-_k)^{-1} + \mathcal{H}^T R^{-1}_k \mathcal{H}]^{-1} = \mathcal{A}_k P^-_k \mathcal{A}^T_k, \tag{13}$$

$$\mu[\hat{z}^+_k] = P^+_k[(P^-_k)^{-1}\mu[\hat{z}^-_k] + \mathcal{H}^T R^{-1}_k y_k], \tag{14}$$

$$\hat{z}^+_{k,i} = \mathcal{A}^T_k(\hat{z}^-_{k,i} - \mu[\hat{z}^-_k]) + \mu[\hat{z}^+_k], \tag{15}$$

where $\hat{y}_k \in \mathbb{R}^m$, $\hat{z}_k \in \mathbb{R}^{n+m}$, $\mathcal{F}_k$ is obtained from the singular value decomposition (SVD) $P^-_k = \mathcal{F}_k D_k \mathcal{F}^T_k$, $\mathcal{G}_k$ is a square root of D_k, U_k is

obtained from the SVD $\mathcal{G}^T_k \mathcal{F}^T_k \mathcal{H}^T R^{-1} \mathcal{H} \mathcal{F}_k \mathcal{G}_k = U_k J_k U^T_k$, B_k is a square root of $(I + J_k)^{-1}$, and $\mathcal{H} = [0_{m \times n} \; I_{m \times m}]$.

Note that (10) assumes no model uncertainty, that is, it assumes no mismatch between the model used to generate the truth data (truth model) and the model used to propagate ensemble members (ensemble model). Additionally, the ensemble model assumes no dynamics for the driver, that is, the prior driver at the current step $\hat{u}^-_{k,i}$ is equal to the posterior driver at the previous step $\hat{u}^+_{k-1,i}$.

To demonstrate driver estimation, (10)–(15) are applied to the linear system (1) and (2). As in the previous section, the ensemble model is initialized from 20 different initial conditions for both $\hat{s}^+_{1,i}$ and $\hat{u}^+_{1,i}$, which are taken from a zero-mean normal distribution with variance 0.4. Fig. 3 shows that the driver estimates converge to constant values and the ensemble spread converges to zero. The ensemble spread converges to zero due to the fact that the ensemble model assumes no dynamics for the driver. The resulting average RMSPE value of 55.45% is an improvement from the values achieved in the previous section by about 7%. Ensemble spread converging to zero and ensemble mean converging to a fixed value are referred to as filter divergence.

To remedy filter divergence, the prior ensemble variance can be inflated as given by

$$\hat{z}^-_{k,i} = \sqrt{\lambda}(\hat{z}^-_{k,i} - \mu[\hat{z}^-_k]) + \mu[\hat{z}^-_k], \tag{16}$$

where $\lambda > 1$ (Anderson and Anderson, 1999). This directly increases both the ensemble spread every step and makes state and driver estimates more uncertain, thereby allowing them to be updated by a greater increment. Fig. 4 demonstrates that $\lambda = 2.0$ decreases the state RMSPE by about 17–38.63% as compared to the data assimilation without ensemble inflation. The driver estimate has approximately the same period and amplitude as the actual driver, but is phase shifted due to the fact that the state estimation compensates for some amount of driver inconsistency and hence driver estimation does not need to estimate the driver perfectly. This would not be the case if data assimilation was performed by estimating the driver without estimating the states.

Fig. 3. Time evolution of EAKF state and driver estimates for the system (1), (2) using driver estimation, but without ensemble inflation. Due to ensemble model assuming no dynamics for the driver, ensemble spread (light blue) approaches zero. Accordingly, the mean estimates converge to fixed values and do not track the true state and driver. (For interpretation of the references to color in this figure caption, the reader is referred to the web version of this paper.)

Fig. 4. Time evolution of EAKF state and driver estimates for the system (1), (2) when driver estimation and ensemble inflation are used. Using variance inflation of $\lambda = 2.0$ results in an average decrease of 17% in the state tracking error.


Fig. 5. Time evolution of EAKF state and driver estimates for the system (1), (2) using the ensemble model (17) where driver estimation and ensemble inflation are used. With ensemble inflated using $\lambda = 2.0$, the driver estimate compensates for model uncertainty and allows minor degradation in state estimate (1%) as compared to the perfect model case.

2.4. EAKF: data assimilation with driver estimation, ensemble inflation and model uncertainty

We now consider the case where the ensemble model (10) is different from the truth model (1) and (2). For example, consider the ensemble model

$$\hat{z}_{k,i}^{-} = [0.1\hat{s}_{k-1,i}^{+} + \hat{u}_{k-1,i}^{+}; \hat{u}_{k-1,i}^{+}; \hat{s}_{k,i}^{+}], \tag{17}$$

where the only difference between this ensemble model and the truth model (1) is in the state dynamics coefficient, which is 0.1 for the ensemble model and 0.5 for the truth model. Since the truth model and the ensemble model are different, the driver estimate has to correct not only for the uncertainty in the measurement y_k, but also for the model uncertainty (Godinez and Koller, 2012). This is often the case with numerical models – they are not perfect representations of nature (Chatfield, 2006), and often have uncertain parameters that are difficult to specify (Pawlowski and Ridley, 2009). Therefore, the uncertainty in the parameters can be remedied by data assimilation. As a result, $\hat{u}_{k-1,i}^{+}$ might not converge to the true driver value, but instead to a value that can drive the state of the ensemble model to the true state. Ensemble inflation has more influence in the uncertain model case as compared to the accurate model case since the driver has to compensate for model uncertainty in addition to the measurement uncertainty. Accordingly, using $\lambda = 2.0$ for the ensemble model (17) to assimilate measurements from the truth model (1), (2), results in the performance shown in Fig. 5. The value of λ is chosen as a compromise between ensemble spread and state tracking. Increasing λ increases the ensemble spread, which in turn allows for large instantaneous change in the estimates. On one hand, a large spread is undesirable since it implies large uncertainty in the estimates, whereas, on the other hand, a small spread slows down the updating of the estimates, and may result in filter divergence.

Fig. 5 demonstrates a 1% degradation in the state RMSPE (with a new value of 39.59%) when compared to the accurate model case. The driver estimate had to deviate farther from true value than in the previous section, but was able to compensate for model uncertainty as shown by relatively small degradation in state RMSPE when compared to values in the previous section.

2.5. EAKF: localization

The previous section considered only a single state and a single driver. Within GITM there are $72 \times 36 \times 50$ grid points with 35 state variables at each grid point (densities, temperatures, and velocities for multiple species), all of which are estimated. Using the equations derived in the previous section for problems of this size is not only prohibitively expensive computationally, but also can produce physically meaningless solutions arising from correlations between distant states and measurements. For example, rapid variations in the temperature above the north pole will have little connection to rapid variations in the temperature above the south pole. This is not to say that the variations at the south pole cannot affect the state variable at the north pole over long period of time, but rather that this effect is not instantaneous. If these correlations are retained, the state estimates could be driven in the wrong direction. For example, if rapid variations in the temperature at the north and south pole were anti-correlated, or phase-shifted, then the cross correlation between the two points could be negative, implying that when a heating event occurred at the north pole, the technique would try to force the south pole to decrease in temperature, which is not physical. Accordingly, the influence of the measurements on the state variables at the current step is spatially confined. This region is defined as a smooth function of the distance between the measurement and the state variable. This localization function takes values between 0 and 1, and multiplies the correlation between the measurement and the state variable (Anderson and Collins, 2007). A value of 1 corresponds to direct connection between the measurement and the state, whereas a value of 0 corresponds to no connection. The values between 0 and 1 vary smoothly as a function of distance. The piecewise continuous function (4.10) in Gaspari and Cohn (1999) is used with a half width of $30°$ in the horizontal direction and 100 km in altitude for all experiments in this paper. The resulting localization region is shown in Fig. 6 for a measurement above Ann Arbor, MI, where the intensity of red indicates the value of the localization function.

In this paper, localization is used for all GITM variables except the driver estimate $\hat{F}_{10.7}$, which is affected by all measurements.

This implies that, for example, the measurements on the night side of the Earth have full impact on the driver.

Although data assimilation can be performed at every step in simple models, this approach becomes impractical in the DART-GITM interface due to the cost of stopping and restarting GITM. Accordingly, even though GITM uses 2 s time step and CHAMP data are available about every 47 s, the assimilation window (time between assimilation steps, t_{aw}) was chosen to be 30 min, which is a compromise between runtime and performance. Smaller assimilation windows were able to improve bias reduction, but required impractically amounts of runtime.

One current feature of DART-GITM interface is that it does not interpolate measurements in time. In other words, all the measurements from current time minus half the assimilation window to current time plus half the assimilation window are used without modification as if they occurred at the current time.

3. Results

All of the data assimilation experiments performed on GITM were run for a period of 2 days during a geomagnetically calm period (December 1–2, 2002). For the simulated measurement
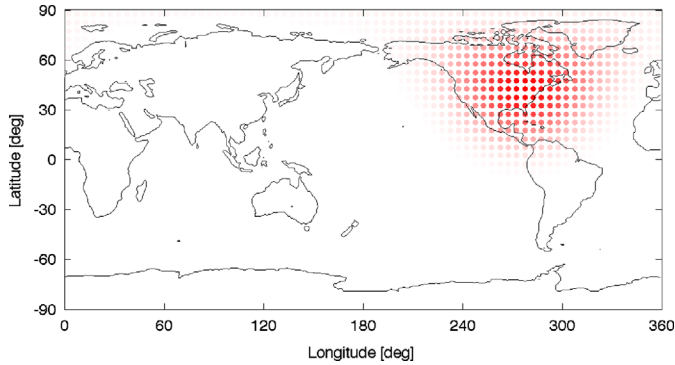
Fig. 6. The localization region for the measurement located above Ann Arbor, MI. Markers are placed at GITM cell center locations. The dark red markers correspond to measurements having a direct effect on the state variable, whereas transparent markers correspond to measurements that have no effect on the state variable. The state variables affected by the measurement above Ann Arbor lie in a projection of a small circle region centered in Ann Arbor. (For interpretation of the references to color in this figure caption, the reader is referred to the web version of this paper.)

experiments the measurement error variance R_k was set to a constant value of 2.6×10^{-13} kg m^{-3}, which was the average CHAMP measurement error for the simulated dates. For the real CHAMP measurement experiment, R_k was set to the measurement error values provided in the CHAMP data files. In this section, the driver estimate was inflated to have a constant variance; the state estimate covariance was not inflated. More precisely, the driver estimate was inflated before computing (11) using

$$\hat{F}_{k,i}^{-} = \frac{\sqrt{\sigma_i^2}}{\sqrt{\sigma^2[\hat{F}_k^{-}]}} \left(\hat{F}_{k,i}^{-} - \mu\left[\hat{F}_k^{-}\right] \right) + \mu\left[\hat{F}_k^{-}\right] \qquad (18)$$

with $\sigma_i^2 = 49$ SFU^2. Keeping the driver variance constant was found to outperform state inflation since the driver estimate variance in the state inflation case would take multiple days to grow to the desired level. On the other hand, driver-only inflation reaches these levels immediately and keeps the variance of the driver estimate constant. While the use of the full state inflation in conjunction with constant drive inflation may produce better results, it was left for future work.

While the EAKF assumes that the state estimates are normally distributed, temperatures, densities and $F_{10.7}$ cannot take on negative values and hence cannot be normal. Accordingly, when estimates of these state variables become negative, their values are set at half the value at the previous step. The initial $\hat{F}_{10.7}$ ensemble distribution was chosen to be normal with a mean of 130 SFU and variance of 25 SFU^2. This distribution created spread in the initial conditions for other variables since all ensemble members were pre-spun for 2 days (11/29–12/01) prior to the start of the data assimilation. The mean was chosen to be below the true value to avoid inadvertently compensating for GITM-CHAMP ρ mismatch. More precisely, since the true $F_{10.7}$ was about 142 SFU during December 1–2, 2002 and GITM, with bias added, underestimated ρ at CHAMP locations, choosing an initial mean $\hat{F}_{10.7}$ to be about 220 SFU would fix the bias. Instead, the mean of 130 SFU was chosen to provide for a more challenging setup. The variance of the initial $\hat{F}_{10.7}$ distribution was chosen as to create the initial variance in ρ to be greater than R_k in the beginning of the assimilation to give the measurements more weight (see Section 2.2 for related discussion).

Also, we note that only $F_{10.7}$ was estimated since it has greater influence on the neutral density during the geomagnetically calm period as opposed to the geomagnetic activity drivers Kp and Ap.
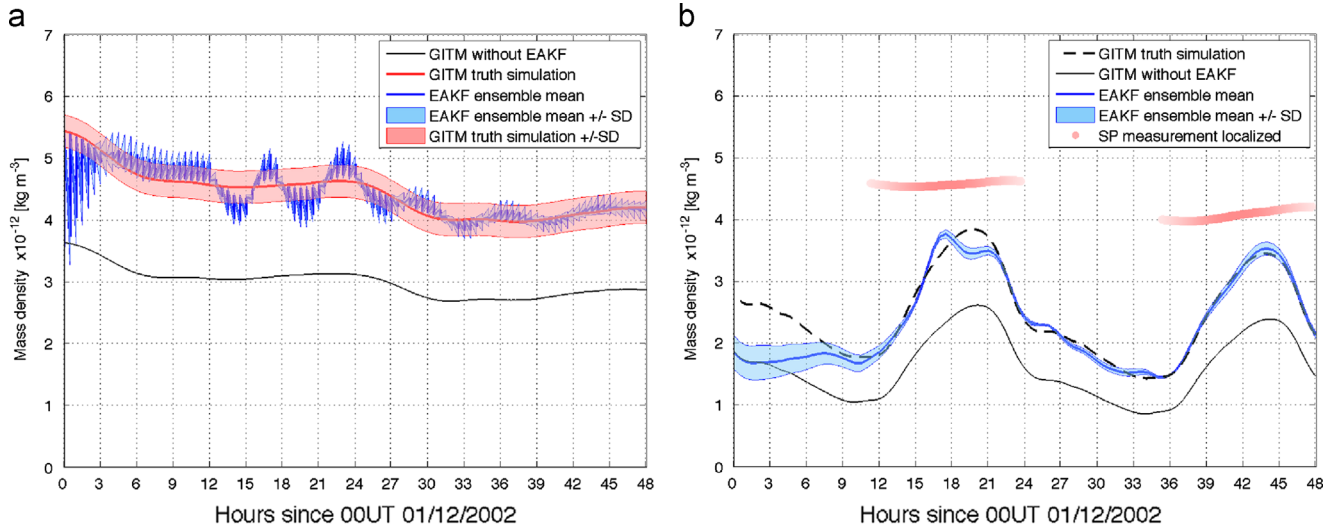
Fig. 7. Time evolution of the simulated (red) and estimated (blue) mass density above the subsolar point (measurement location) in (a) and at the 400 km gridpoint closest to Ann Arbor, MI in (b). Estimated mass density above the subsolar point converged to the true value within 3 h, whereas density at the Ann Arbor gridpoint needed 12 h to converge. (For interpretation of the references to color in this figure caption, the reader is referred to the web version of this paper.)

Performing the experiments described in this section during the geomagnetically disturbed periods is the subject of future work and will likely benefit from estimation of the geomagnetic indices and, possibly, the photoelectron heating efficiency.

Finally, the number of ensemble members was fixed at 20 for this study due to the initial limit on computational resources. The effect of changing the number of ensemble members was not investigated in this study. More precisely, $5° \times 5°$ GITM runs on 32 processors, which for 20 instances entails 640 processors. To facilitate the simulations, NASA's Pleiades supercomputer was used. Investigating the effect of the number of ensemble members is the subject of future work.

3.1. Simulated data above the subsolar point

As a first example, a simulated measurement of the density ρ_{sp} at about 400 km above the subsolar point is obtained from a GITM truth simulation. Since it is impractical to measure ρ_{sp} in practice, this example illustrates EAKF in the case where the measurement is directly linked to the quantity to be estimated. We thus record ρ_{sp} at 1-min intervals from the truth model for use in the EAKF as measurements. Truth GITM is simulated for 2 days (December 1–2,


Fig. 9. Time evolution of $F_{10.7}$ for the truth simulation (dashed black), nominal no data assimilation simulation (solid black), and the ensemble mean (dark blue). The ensemble mean oscillated around the true value and converged to the true value by the end of the simulation. (For interpretation of the references to color in this figure caption, the reader is referred to the web version of this paper.)
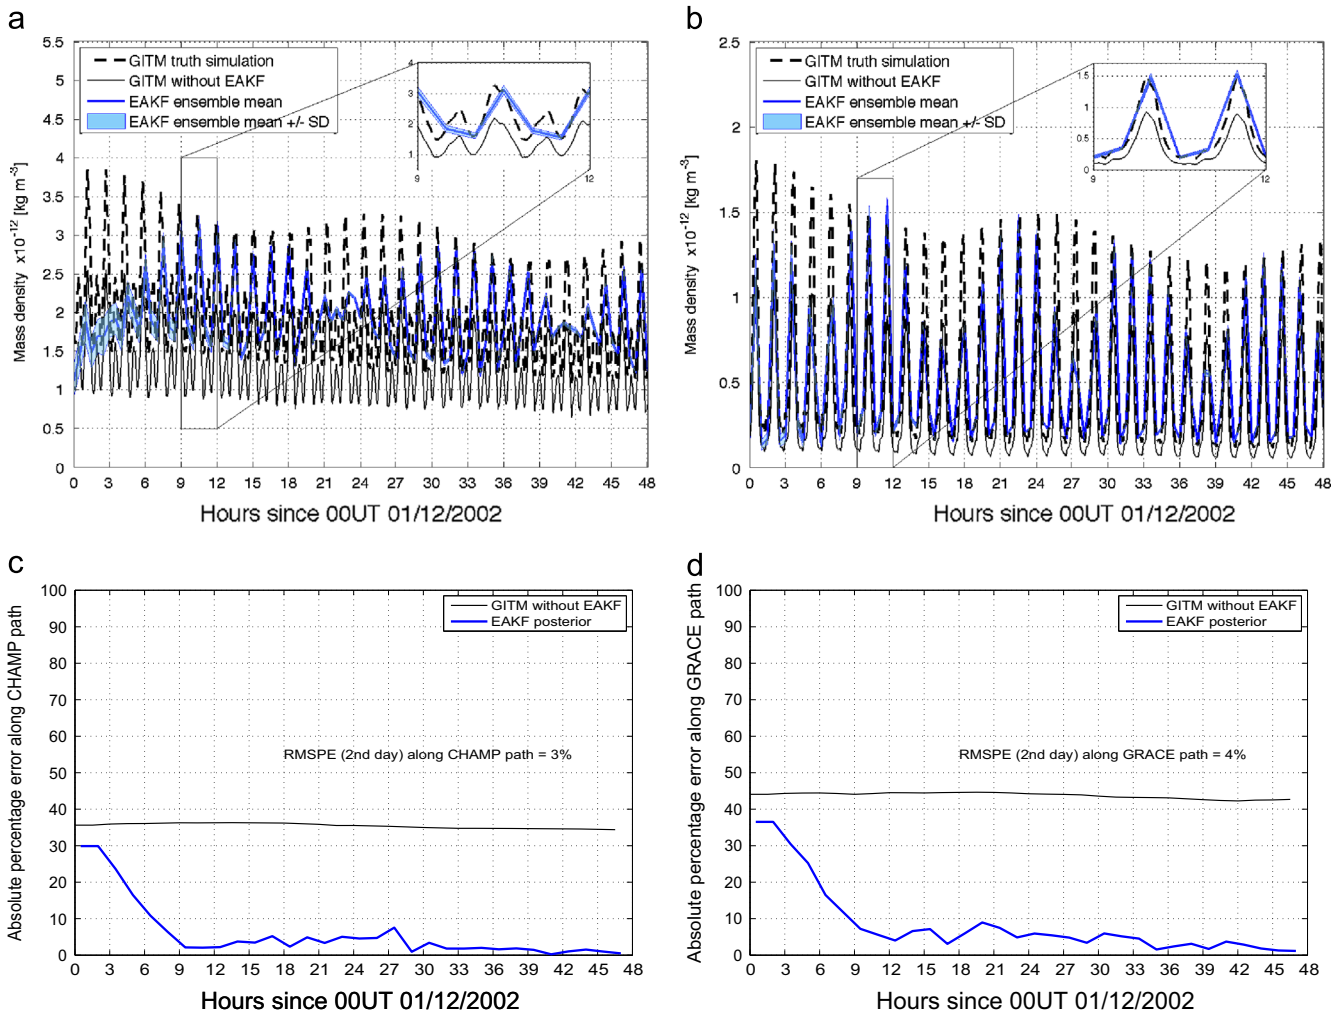

Fig. 8. (a) shows that the time evolution of mass density at CHAMP locations demonstrates that ρ converged to the true value within 9 h when measurements were taken at the subsolar point. (b) shows mass density estimate at GRACE locations and demonstrates the same rate of convergence. (c) and (d) reinterpret (a) and (b) by plotting the orbit averages of the errors between simulated and estimated data. This figure demonstrates that estimated mass density converged to true mass density at locations other than the measurement location (subsolar point).

2002) with the true $F_{10.7}$ (mean value of 142 SFU). With data assimilation performed every 30 min, Fig. 7 shows results from the truth model and ensemble estimates at two locations, namely, the subsolar point (measurement location) and the GITM grid point closest to Ann Arbor, MI (diagnostics location). Fig. 7 (b) demonstrates that state estimates (blue) at the subsolar point converge to the true state (red). The measurement uncertainty used in this experiment is the average value of CHAMP measurement error data (standard deviation) for the December 1–3, 2002 and is shown in light red. The solid black line shows the estimated density at the subsolar point without data assimilation with $F_{10.7} = 130$ SFU (the mean of the initial $\hat{F}_{10.7}$ distribution). The assimilated mean deviates from this passive trajectory and converges to approximately the true state within 3 h.

Fig. 7(b) shows that even at a location different from the measurement location (the GITM grid point closest to Ann Arbor, MI) the mean state estimate (blue, $\hat{\rho}_{aa}$) converged to the true state (dashed black, ρ_{aa}). At this location the convergence is slower since Ann Arbor only entered the subsolar point localization region between 12 and 24 UT each day.

Fig. 8(a) and (b) shows mass density estimates at CHAMP and GRACE locations ($\hat{\rho}_{ch}$ and $\hat{\rho}_{gr}$, respectively). The inset in (a) demonstrates that by about 09 UT on December 1, 2002 the CHAMP ensemble mean $\mu[\hat{\rho}_{ch}]$ converged to the true value ρ_{ch}. Similarly, (a) leads us to conclude that convergence took place along the GRACE trajectory as well. The absolute percentage error (APE) values are computed for (a) and (b), as

$$\text{APE}_k^- = 100 \times \frac{|\mu[\hat{\rho}_{ch,k,i}^-] - \rho_{ch,k}|}{|\rho_{ch,k}|}, \tag{19}$$

and are shown in (c) and (d). These plots demonstrate that the subsolar point measurement not only removed the bias at the subsolar point, but also at the CHAMP and GRACE locations. More precisely, the RMSPE^- at the CHAMP locations over the second day decreased from 36% to 3% when compared to the bias in the nominal case (the truth model with $F_{10.7} = 142$ SFU and the no data assimilation GITM using $F_{10.7} = 130$ SFU). Similarly, the RMSPE^- at the GRACE location decreased from 43% to 4%. Finally, Fig. 9 shows that the $F_{10.7}$ estimate converged to the true value by the end of the simulated period.
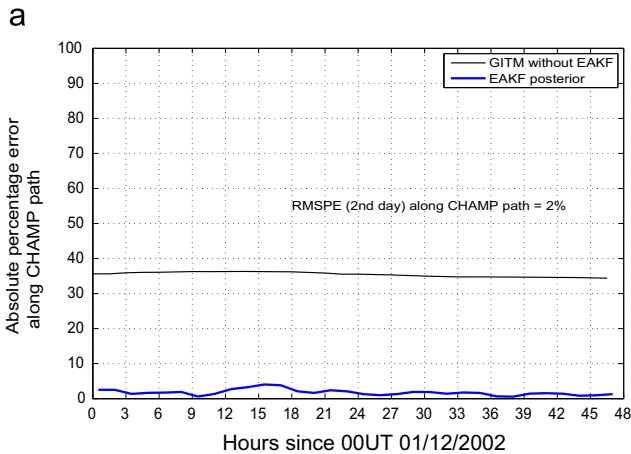
3.2. Simulated CHAMP data

As a more physically realistic case, measurements at the CHAMP locations are used. Fig. 1 shows that the CHAMP orbit does not pass through the subsolar point, and hence this case is more challenging than the previous example. Fig. 10 demonstrates that even in this seemingly harder case, EAKF was able to decrease the prior RMSPE from 36% to 2% along CHAMP trajectory and from 43% to 4% along GRACE trajectory compared to the no data assimilation case. The values of $\hat{F}_{10.7}$ shown in Fig. 11 demonstrate that the technique initially overestimated the true value, but then converged close to the true value within 15 h. Indeed, the convergence was significantly faster than in the previous example.

3.3. Real CHAMP data

For the final case, the mass density and uncertainty estimated by CHAMP are used to attempted to correct a GITM simulation with bias added. Because the true state was unknown, data from


Fig. 11. Time evolution of $F_{10.7}$ for the true, nominal and ensemble simulations. The ensemble mean estimate initially overestimates the true value, but then converges within 15 h.
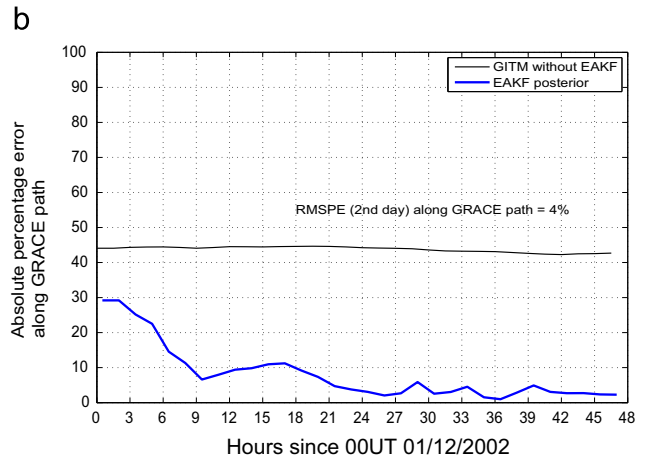

Fig. 10. Orbit averaged mass density absolute percentage errors at (a) CHAMP and (b) GRACE locations. Simulated CHAMP data were used as measurements in the EAKF, whereas the simulated GRACE values were only used for validation. Density estimate at CHAMP locations dropped immediately to about 3%, whereas the estimate at GRACE locations decreased only after about 9 h due to altitude difference between these satellites. The bias in densities at both locations was decreased and the final error level was lower at the measurement location (CHAMP) than at the validation location (GRACE).

GRACE were used as the validation metric. Since GITM with proper high-latitude driving has a relatively low bias compared to CHAMP data, GITM simulations were conducted with constant high-latitude driving, to intentionally introduce bias. Fig. 12 shows that the bias at the CHAMP locations with data assimilation and driver estimation is reduced from 73% to 7%. At the GRACE satellite location, however, the bias was only marginally reduced. Even though the GRACE RMSPE is reduced from 76% to 52%, the large remaining bias suggests that either more than one driver needs to be estimated in order to remove bias at multiple locations or that more data needs to be assimilated to remove the bias over the whole thermosphere. Also, GRACE is at a higher altitude and different local time sector than CHAMP. This means that the assimilation localization region surrounding CHAMP may rarely encounter the GRACE satellite. Furthermore, the difference in the amount of bias removed between the satellites may say something about the method in which the bias is removed. In this case, more heat on the day side is added thereby making the thermosphere warmer, but not in the same way as increased high-latitude heating would. Therefore, while the simulation results improve at the GRACE locations, they are not dramatically improved,
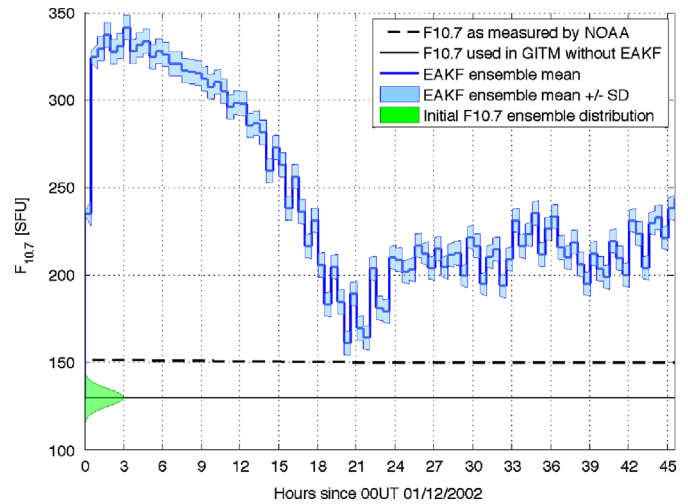

Fig. 13. Time evolution of $F_{10.7}$ for the true, nominal and ensemble simulations. $F_{10.7}$ estimate took on higher values than what was measured by NOAA, in order to drive GITM closer to the CHAMP data.
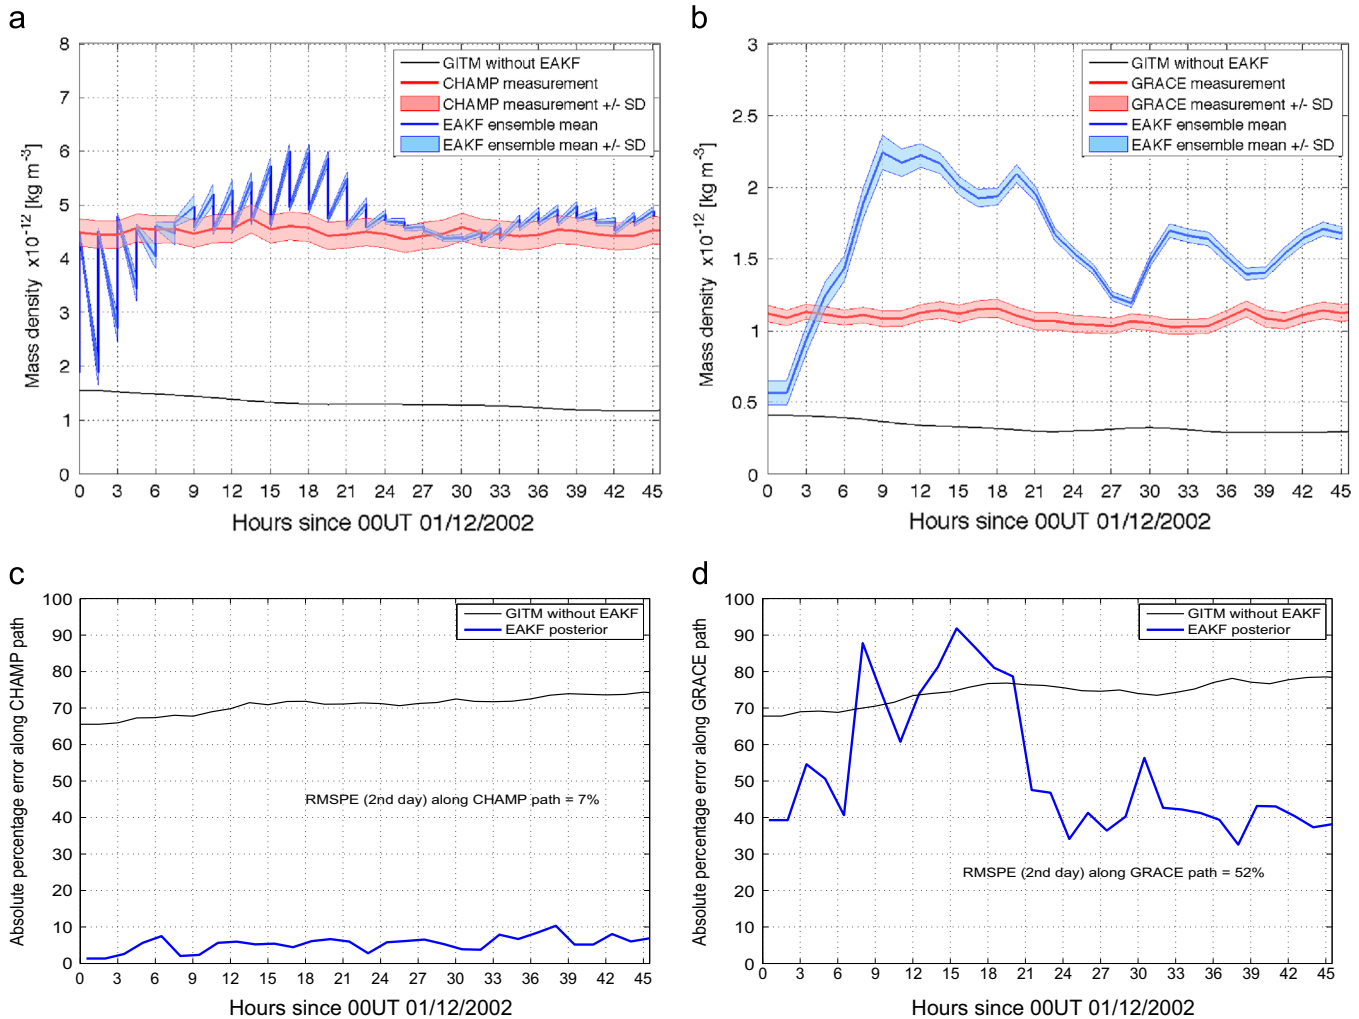

Fig. 12. Orbit averaged mass density measurements at CHAMP and GRACE are shown in (a) and (b) respectively, and absolute percentage errors are shown in (c) and (d) respectively. Real CHAMP data were used as measurements in the EAKF, whereas GRACE data were used only for validation. The bias in density at CHAMP locations was decreased over the second day from 57% to 7%, whereas the bias at GRACE location was only decreased by a smaller amount (from 70% to 52%, see text for relevant discussion).

pointing to the possible need to remove the bias in a more physically correct way. Fig. 13 shows the $F_{10.7}$ values used to achieve the aforementioned performances and demonstrates that the driver had to take on higher values (about 220 SFU instead of nominal 142 SFU) to compensate for the model bias. If this bias was removed in a more physically meaningful way, such as allowing the high-latitude dynamics of behave as they should, the bias compensation may be better and will be explored in further studies.

4. Discussion and conclusions

This study described the preliminary implementation of the DART-GITM interface. In particular, it demonstrated how both simulated and real CHAMP ρ measurements can be assimilated into GITM. In addition to estimating GITM states such as densities, temperatures, and velocities, this interface estimated one of GITM drivers, namely, $F_{10.7}$. This driver was not estimated for the purposes of knowing $F_{10.7}$ more precisely, but instead for the purposes of driving GITM's estimate of ρ closer to CHAMP measurements. Accordingly, it was found that in the simulated CHAMP measurements case, the good estimate of $F_{10.7}$ was able to drive GITM to the point of decreasing the bias in the simulated mass densities at the CHAMP locations from 36% to 2% and from 43% to 4% at the validation locations (GRACE orbit). This technique produced greater decrease when real CHAMP data were used as measurements in EAKF and GRACE data were used for validation. More precisely, assimilating real CHAMP data and estimating $F_{10.7}$ in a biased GITM decreased the mass density bias from 58% to 7% for CHAMP locations and from 77% to 52% for GRACE locations when compared to not performing data assimilation.

The relatively large final GRACE bias (52% compared to CHAMP's 7%) could possibly be explained by the way these data were derived. In particular, Matsuo et al. (2012) mentioned that CHAMP and GRACE data were found to be inconsistent when compared to the CTIPe outputs and attributes this disparity to the "uncertainty of the drag coefficient assumptions employed in the retrieval as well as the different orbital altitudes." The authors scaled the GRACE densities by an arbitrary number to remove the apparent inconsistency. This solution might have reduced the observed GRACE RMSPE in this study, but is left for future work.

These preliminary results can be further improved by tuning parameters and relaxing of some assumptions. In particular, the interface parameters that require tuning include the number of ensemble members, the data assimilation temporal window, and observational error variance R. Initial experiments with the number of ensemble members and the assimilation window show that these parameters only affect the speed of convergence of the state and driver estimates to the true values (not shown), but a more thorough study is needed. The possible interface improvements mentioned above include using log-normal distributions instead of normal distributions for densities and temperatures or at least saturating negative values to zero instead of half the previous value, defining a localization region for the driver, and implementing temporal localization of the measurements. Also, investigating full state inflation in conjunction with constant driver inflation is subject for future work.

Future work will include estimating more drivers, such as the auroral power and cross polar cup potential, and assimilating more measurements, such as the total electron content measurements made by GPS, which will help to extend this study from a thermosphere into the ionosphere. Additionally, this study considered only the orbit averaged estimates of bias and therefore did not explore latitude dependence of the assimilation performance. This research direction might be of interest due to high variability

of the energy sources in the polar region. Lastly, since this study considered only the geomagnetically calm period, it would be interesting to see whether the current setup is capable of producing equally good estimates during a geomagnetic storm.

The main goal in this demonstration was to explore whether it is possible to remove the bias in a model using data assimilation along with driver estimation. The last example demonstrates that at the location where data are ingested, the bias is almost completely removed (due mostly to the data assimilation), while, at other locations, the bias is reduced. This shows that the technique works but might be dependent on choosing the correct driver to estimate.

References

Anderson, J., Collins, N., 2007. Scalable implementations of ensemble filter algorithms for data assimilation. Journal of Atmospheric and Oceanic Technology 24, 1452–1463.

Anderson, J., Hoar, T., Raeder, K., Liu, H., Collins, N., Torn, R., Avellano, A., 2009. The data assimilation research testbed: a community facility. Bulletin of the American Meteorological Society 90, 1283–1296.

Anderson, J.L., 2001. An ensemble adjustment Kalman Filter for data assimilation. Monthly Weather Review 129, 2884–2903.

Anderson, J.L., Anderson, S.L., 1999. A monte carlo implementation of the nonlinear filtering problem to produce ensemble assimilations and forecasts. Monthly Weather Review 127, 2741–2758.

Bilitza, D., 2001. International reference ionosphere 2000. Radio Science 36, 261–275.

Chatfield, C., 2006. Model Uncertainty. John Wiley & Sons, Ltd.

Deng, Y., Richmond, A., Ridley, A., Liu, H., 2008. Assessment of the non-hydrostatic effect on the upper atmosphere using a general circulation model (GCM). Geophysical Research Letters 35, L01104.

Evensen, G., 1994. Sequential data assimilation with a nonlinear quasi-geostrophic model using monte carlo methods to forecast error statistics. Journal of Geophysical Research 99, 10143–10162.

Fuller-Rowell, T., Evans, D., 1987. Height-integrated Pedersen and Hall conductivity patterns inferred from the TIROS-NOAA satellite data. Journal of Geophysical Research: Space Physics (1978–2012) 92, 7606–7618.

Gaspari, G., Cohn, S.E., 1999. Construction of correlation functions in two and three dimensions. Quarterly Journal of the Royal Meteorological Society 125, 723–757.

Godinez, H., Koller, J., 2012. Localized adaptive inflation in ensemble data assimilation for a radiation belt model. Space Weather: The International Journal of Research and Applications 10, S08003.

Hedin, A.E., 1987. MSIS-86 thermospheric model. Journal of Geophysical Research 92, 4649–4662.

Kersley, L., Pryse, S., Denton, M., Bust, G., Fremouw, E., Secan, J., Jakowski, N., Bailey, G., 2005. Radio tomographic imaging of the northern high-latitude ionosphere on a wide geographic scale. Radio science 40, 1–9.

Khade, V., Hansen, J., Reid, J., Westphal, D., 2012. Ensemble filter based estimation of spatially distributed parameters in a mesoscale dust model: experiments with simulated and real data. Atmospheric Chemistry and Physics Discussions 12.

Kim, I., 2008. Large scale data assimilation with application to the ionosphere–thermosphere. ProQuest.

Kim, I., Pawlowski, D., Ridley, A., Bernstein, D., 2008. Localized data assimilation in the ionosphere–thermosphere using a sampled-data unscented Kalman filter. In: American Control Conference, IEEE, 2008, pp. 1849–1854.

Lee, I.T., Matsuo, T., Richmond, A.D., Liu, J.Y., Wang, W., Lin, C.H., Anderson, J.L., Chen, M.Q., 2012. Assimilation of FORMOSAT-3/COSMIC electron density profiles into a coupled thermosphere/ionosphere model using ensemble Kalman filtering. Journal of Geophysical Research 117.

Li, J., Liu, H., 2009. Improved hurricane track and intensity forecast using single field-of-view advanced ir sounding measurements. Geophysical Research Letters 36, L11813.

Mandrake, L., Wilson, B., Wang, C., Hajj, G., Mannucci, A., Pi, X., 2005. A performance evaluation of the operational Jet Propulsion Laboratory/University

of Southern California global assimilation ionospheric model (JPL/USC GAIM). Journal of geophysical research 110, A12306.

Mannucci, A., Wilson, B., Yuan, D., Ho, C., Lindqwister, U., Runge, T., 1998. A global mapping technique for GPS-derived ionospheric total electron content measurements. Radio Science 33, 565–582.

Matsuo, T., Fedrizzi, M., Fuller-Rowell, T.J., Codrescu, M.V., 2012. Data assimilation of thermospheric mass density. Space Weather 10, S05002.

Millward, G., Moffett, R., Quegan, S., Fuller-Rowell, T., 1996. A coupled thermosphere-ionosphere–plasmasphere model (CTIP). In: Handbook of Ionospheric Models, pp. 239–279.

Millward, G., Müller-Wodarg, I., Aylward, A., Fuller-Rowell, T., Richmond, A., Moffett, R., 2001. An investigation into the influence of tidal forcing on F region equatorial vertical ion drift using a global ionosphere–thermosphere model with coupled electrodynamics. Journal of Geophysical Research, A: Space Physics 106, 24.

Pawlowski, D., Ridley, A., 2008. Modeling the thermospheric response to solar flares. Journal of Geophysical Research 113, A10309.

Pawlowski, D., Ridley, A., 2009. Quantifying the effect of thermospheric parameterization in a global model. Journal of Atmospheric and Solar-Terrestrial Physics 71, 2017–2026.

Perbos, J., Escudier, P., Parisot, F., Zaouche, G., Vincent, P., Menard, Y., Manon, F., Kunstmann, G., Royer, D., Fu, L., 2003. Jason-1: assessment of the system performances special issue: Jason-1 calibration/validation. Marine Geodesy 26, 147–157.

Pi, X., Wang, C., Hajj, G.A., Rosen, G., Wilson, B.D., Bailey, G.J., 2003. Estimation of $E \times B$ drift using a global assimilative ionospheric model: an observation system simulation experiment. Journal of Geophysical Research 108, 1075.

Reigber, C., Lühr, H., Schwintzer, P., 2002. CHAMP mission status. Advances in Space Research 30, 129–134.

Ridley, A., Deng, Y., Tóth, G., 2006. The global ionosphere–thermosphere model. Journal of Atmospheric and Solar-Terrestrial Physics 68, 839–864.

Russo, G., Di Bernardo, M., Sontag, E., 2010. Global entrainment of transcriptional systems to periodic inputs. PLoS computational biology 6, e1000739.

Scherliess, L., Schunk, R., Sojka, J., Thompson, D., 2004. Development of a physics-based reduced state Kalman filter for the ionosphere. Radio Science 39, RS1S04.

Schunk, R., Scherliess, L., Sojka, J., Thompson, D., Anderson, D., Codrescu, M., Minter, C., Fuller-Rowell, T., Heelis, R., Hairston, M., et al., 2004. Global assimilation of ionospheric measurements (GAIM). Radio Science 39, RS1S02.

Sontag, E., 2010. Contractive systems with inputs. Perspectives in Mathematical System Theory, Control, and Signal Processing, 217–228.

Stensrud, D., Gao, J., 2010. Importance of horizontally inhomogeneous environmental initial conditions to ensemble storm-scale radar data assimilation and very short-range forecasts. Monthly Weather Review 138, 1250–1272.

Storz, M., Bowman, B., Branson, M., Casali, S., Tobiska, W., 2005. High accuracy satellite drag model (HASDM). Advances in Space Research 36, 2497–2505.

Sutton, E., 2009. Normalized force coefficients for satellites with elongated shapes. Journal of Spacecraft and Rockets 46, 112–116.

Sutton, E., Nerem, R., Forbes, J., 2007. Density and winds in the thermosphere deduced from accelerometer data. Journal of Spacecraft and Rockets 44, 1210–1219.

Tapley, B., Bettadpur, S., Watkins, M., Reigber, C., 2004. The gravity recovery and climate experiment: mission overview and early results. Geophysical Research Letters 31, L09607.

Tarantola, A., 2005. Inverse Problem Theory and Methods for Model Parameter Estimation. Siam, Philadelphia.

Wang, C., Hajj, G., Pi, X., Rosen, I.G., Wilson, B., 2004. Development of the global assimilative ionospheric model. Radio Science 39, RS1S06.

Weimer, D., 1996. A flexible, IMF dependent model of high-latitude electric potentials having "space weather" applications. Geophysical Research Letters 23, 2549–2552.

Yiğit, E., Ridley, A., 2011. Effects of high-latitude thermosphere heating at various scale sizes simulated by a nonhydrostatic global thermosphere–ionosphere model. Journal of Atmospheric and Solar-Terrestrial Physics 73, 592–600.