



Assimilating compact phase space retrievals of atmospheric composition with WRF-Chem/DART: a regional chemical transport/ensemble Kalman filter data assimilation system

Arthur P. Mizzi¹, Avelino F. Arellano Jr.², David P. Edwards¹, Jeffrey L. Anderson³, and Gabriele G. Pfister¹

¹National Center for Atmospheric Research, Atmospheric Chemistry Observation and Modeling Laboratory, Boulder, CO, USA

²University of Arizona, Department of Hydrology and Atmospheric Science, Tucson, AZ, USA

³National Center for Atmospheric Research, Institute for Applied Mathematics, Boulder, CO, USA

Correspondence to: Arthur P. Mizzi (mizzi@ucar.edu)

Received: 3 June 2015 – Published in Geosci. Model Dev. Discuss.: 8 September 2015

Revised: 20 January 2016 – Accepted: 10 February 2016 – Published: 4 March 2016

Abstract. This paper introduces the Weather Research and Forecasting Model with chemistry/Data Assimilation Research Testbed (WRF-Chem/DART) chemical transport forecasting/data assimilation system together with the assimilation of *compact phase space retrievals* of satellite-derived atmospheric composition products. WRF-Chem is a state-of-the-art chemical transport model. DART is a flexible software environment for researching ensemble data assimilation with different assimilation and forecast model options. DART's primary assimilation tool is the ensemble adjustment Kalman filter. WRF-Chem/DART is applied to the assimilation of Terra/Measurement of Pollution in the Troposphere (MOPITT) carbon monoxide (CO) trace gas retrieval profiles. Those CO observations are first assimilated as quasi-optimal retrievals (QORs). Our results show that assimilation of the CO retrievals (i) reduced WRF-Chem's CO bias in retrieval and state space, and (ii) improved the CO forecast skill by reducing the Root Mean Square Error (RMSE) and increasing the Coefficient of Determination (R^2). Those CO forecast improvements were significant at the 95 % level.

Trace gas retrieval data sets contain (i) large amounts of data with limited information content per observation, (ii) error covariance cross-correlations, and (iii) contributions from the retrieval prior profile that should be removed before assimilation. Those characteristics present challenges to the assimilation of retrievals. This paper addresses those challenges by introducing the assimilation of compact phase

space retrievals (CPSRs). CPSRs are obtained by preprocessing retrieval data sets with an algorithm that (i) compresses the retrieval data, (ii) diagonalizes the error covariance, and (iii) removes the retrieval prior profile contribution. Most modern ensemble assimilation algorithms can efficiently assimilate CPSRs. Our results show that assimilation of MOPITT CO CPSRs reduced the number of observations (and assimilation computation costs) by $\sim 35\%$, while providing CO forecast improvements comparable to or better than with the assimilation of MOPITT CO QORs.

1 Introduction

There is increased international interest in chemical weather forecasting (Kukkonen et al., 2012; MACC-II Final Report, 2014). Such forecasts rely on coupled forecast model–data assimilation systems that ingest a combination of remotely sensed and in situ atmospheric composition observations together with conventional meteorological observations. Generally the remotely sensed observations come in the form of trace gas retrievals. Examples include carbon monoxide (CO) total and partial column or profile retrievals from the Terra/Measurement of Pollution in the Troposphere (MOPITT) and Metop/Infrared Atmospheric Sounding Interferometer (IASI) instruments. The associated data sets are characterized by large numbers of observations with limited information per observation. Such remotely sensed data have

been assimilated in various settings (e.g. Bei et al., 2008; Herron-Thorpe et al., 2012; Klonecki et al., 2012; Gaubert et al., 2014), but there have been only a few papers addressing data compression strategies. Two such papers were Joiner and da Silva (1998) and Migliorini et al. (2008). This article is inspired by their research and introduces an efficient assimilation strategy that reduces the number of MOPITT CO retrieval observations by $\sim 35\%$. Greater reductions are possible, for example, with IASI CO and ozone (O_3) retrievals, and depend on the number of (i) levels in the retrieval profile and (ii) linearly independent pieces of information in the retrieval profile.

Joiner and da Silva (1998) first proposed the idea of using information content to reduce the number of retrieval observations. They suggested projecting retrievals onto the eigenvectors of the observation error covariance matrix and zeroing those coefficients with little or no information. Their approach evolved from one-dimensional retrieval algorithms (e.g. Twomey, 1974; Smith and Woolf, 1976; Thompson, 1992).

Retrievals are often obtained using the optimal estimation method of Rodgers (2000) to obtain solutions to the *retrieval equation*

$$\mathbf{y}_r = \mathbf{A}\mathbf{y}_t + (\mathbf{I} - \mathbf{A})\mathbf{y}_a + \boldsymbol{\varepsilon}, \quad (1)$$

where \mathbf{y}_r is the retrieval profile, \mathbf{A} is the averaging kernel, \mathbf{y}_t is the true atmospheric profile (unknown), \mathbf{I} is the identity matrix, \mathbf{y}_a is the retrieval prior profile, and $\boldsymbol{\varepsilon}$ is the measurement error in retrieval space with error covariance \mathbf{E}_m – the measurement error covariance in retrieval space. Joiner and da Silva (1998) proposed projecting Eq. (1) onto the trailing left singular vectors from the singular value decomposition (SVD) of (i) $(\mathbf{I} - \mathbf{A})$ (their method 1) or (ii) the smoothing error $\mathbf{E}_s = (\mathbf{A} - \mathbf{I})\mathbf{P}_p(\mathbf{A} - \mathbf{I})^T$, where \mathbf{P}_p is the retrieval prior error covariance (their method 2). Those projections removed the components of the retrieval that the instrument could not measure with sufficient sensitivity. They called that approach *null-space filtering*.

Joiner and da Silva (1998) recognized that when the retrieval was strongly constrained by the retrieval prior profile, the assumptions underlying null-space filtering were invalid. For such retrievals, they proposed filtering based on an eigen-decomposition of \mathbf{E}_m (their PED (partial eigen-decomposition) retrievals). Their analysis showed that PED retrievals were well conditioned and independent of the retrieval prior profile.

Migliorini et al. (2008) noted that the Joiner and DaSilva (1998) filtering depended on their truncation criteria and was therefore somewhat arbitrary. They also showed it was possible to achieve similar filtering results with an alternative approach that used a more well-defined truncation criterion. Migliorini et al. (2008) rearranged Eq. (1) to obtain

$$\mathbf{y}_r - (\mathbf{I} - \mathbf{A})\mathbf{y}_a = \mathbf{A}\mathbf{y}_t + \boldsymbol{\varepsilon}. \quad (2)$$

Following their terminology, we call the left side of Eq. (2) a *quasi-optimal* retrieval (QOR). Migliorini et al. (2008) noted that \mathbf{E}_m was unlikely to be diagonal and likely to be poorly conditioned. To address those issues, they applied a SVD transform to Eq. (2) based on the leading left singular vectors of \mathbf{E}_m (similar to that proposed by Anderson, 2003) – this step provided diagonalization of \mathbf{E}_m and used a well-determined truncation cutoff. They also applied a scaling based on the inverse square root of the associated singular values – this step improved numerical conditioning.

Migliorini et al. (2008) continued to reduce the dimension of $\mathbf{A}\mathbf{y}_t$ (i.e., the number of observations) by neglecting those elements whose variability was smaller than the measurement error standard deviation (unity in their rotated and scaled system). They proposed identifying those elements with an eigen-decomposition of the covariance of $\mathbf{A}\mathbf{y}_t$. Since that covariance is generally unknown, they replaced it with the forecast error covariance and showed that the resulting dimension was approximately equal to the number of independent linear functions that could be measured to better than noise level.

A more recent paper (Migliorini 2012) shows that retrievals can be transformed to represent only the portion of the state that is well constrained by the original radiance measurements when two requirements are satisfied: (i) the radiance observation operator is approximately linear in a region of state space centered on the retrieval and with a radius on the order of the retrieval error, and (ii) the prior information used to constrain the retrieval does not underrepresent the variability of the state. Migliorini (2012) proves that when those conditions are met the assimilation of radiances is equivalent to the assimilation of retrievals. The Migliorini (2012) analysis shows that it is possible to use information from the retrieval algorithm to compress information in the transformed retrievals

In this paper, we propose an approach that achieves results similar to (i) Migliorini et al. (2008) without needing to approximate the covariance of $\mathbf{A}\mathbf{y}_t$, and (ii) Migliorini (2012) without needing information about the retrieval algorithm. Our goal is to compress the retrievals and remove those components that are not dependent on the measurements. In so doing we expect to make the assimilation of retrievals more computationally efficient. The rest of this paper is organized as follows: Sect. 2 introduces compact phase space retrievals, and Sect. 3 introduces the WRF-Chem/DART regional chemical weather forecast/data assimilation system. Section 4 discusses our experimental design including the study period, model domain, initial/boundary conditions, and relevant WRF-Chem/DART parameter settings. Section 5 discusses the observations that were assimilated in the various experiments described in Sect. 6. The results from those experiments are presented in Sect. 7, and we end with a summary of our thoughts and conclusions in Sect. 8.

2 Assimilation of compact phase space retrievals

We can rewrite Eq. (2) as

$$\mathbf{y}_r - (\mathbf{I} - \mathbf{A})\mathbf{y}_a - \boldsymbol{\varepsilon} = \mathbf{A}\mathbf{y}_t. \quad (3)$$

In Eq. (3) the averaging kernel \mathbf{A} is singular and has low rank. Therefore the information content for each component of $\mathbf{A}\mathbf{y}_t$ is relatively small, and the assimilation is inefficient because one must assimilate the entire profile to get the same information as can be compressed into a number of processed observations equal to the rank of \mathbf{A} . To compress Eq. (3), we propose transforming it with the leading left singular vectors from a SVD of the averaging kernel; i.e., $\mathbf{A} = \mathbf{U}\mathbf{S}\mathbf{V}^T$. We denote the truncated system with the subscript zero so that $\mathbf{A}_0 = \mathbf{U}_0\mathbf{S}_0\mathbf{V}_0^T$; the truncated averaging kernel is obtained by setting the trailing singular values (i.e., singular values that were less than 1.0×10^{-4}) and vectors to zero. The transformed system has the form

$$\mathbf{U}_0^T(\mathbf{y}_r - (\mathbf{I} - \mathbf{A})\mathbf{y}_a - \boldsymbol{\varepsilon}) = \mathbf{S}_0\mathbf{V}_0^T\mathbf{y}_t. \quad (4)$$

Migliorini et al. (2008) showed that subtracting $(\mathbf{I} - \mathbf{A})\mathbf{y}_a$ from Eq. (3) removes all contribution from the retrieval prior profile. Equations (3) and (4) confirm their result because the leading left singular vectors of \mathbf{A} span its range, so the left side of Eq. (3) should project completely onto \mathbf{U}_0^T . Following that transform \mathbf{E}_m becomes $\mathbf{U}_0^T\mathbf{E}_m\mathbf{U}_0$, which may still be non-diagonal and poorly conditioned. Therefore, we apply an SVD transform and inverse scaling similar to that used by Migliorini et al. (2008). If the SVD of $\mathbf{U}_0^T\mathbf{E}_m\mathbf{U}_0$ has the form $\mathbf{U}_0^T\mathbf{E}_m\mathbf{U}_0 = \Phi\Sigma\Psi^T$, the transformed and conditioned form of Eq. (4) is

$$\Sigma^{-1/2}\Phi^T\mathbf{U}_0^T(\mathbf{y}_r - (\mathbf{I} - \mathbf{A})\mathbf{y}_a - \boldsymbol{\varepsilon}) = \Sigma^{-1/2}\Phi^T\mathbf{S}_0\mathbf{V}_0^T\mathbf{y}_t. \quad (5)$$

Our approach compresses Eq. (3) so that the dimension of the compact phase space retrieval (CPSR) profile on the left side of Eq. (5) is identical to the number of independent linear functions of the atmospheric profile to which the instrument is sensitive. This method is different from that of Migliorini et al. (2008) because it compresses the quasi-optimal retrieval observations based on a linear independence analysis and relies on the assimilation system to decide how much weight to give the observations. The approach of Migliorini et al. (2008) reduces the number of observations based on an uncertainty analysis independent of the assimilation system. Our approach identifies all linearly independent information contained in the QOR profile (through projection of the QOR profile onto the left non-zero singular vectors of the averaging kernel). The approach of Migliorini et al. (2008) may (i) discard some linearly independent information because the left non-zero singular vectors of the observation error covariance are not necessarily a basis for the space of QORs, and (ii) discard some linearly independent information through their uncertainty analysis. Finally, our

approach relies on two transforms: (i) a compression transform (based on the left non-zero singular vectors of the averaging kernel), and (ii) a diagonalization transform (based on the left non-zero singular vectors of the compressed observation error covariance). The approach of Migliorini et al. (2008) uses two diagonalization transforms – the first based on the observation error covariance and the second based on the transformed forecast error covariance in observation space. Our diagonalization transform is analogous to their first diagonalization transform except we apply it to the compressed observation error covariance, and they apply it to the *untransformed* observation error covariance. As in Migliorini et al. (2008), the final form of our observation error covariance is the truncated identity matrix.

The assimilation of CPSRs should produce results similar to the assimilation of QORs except for (i) the effect of assimilation sub-processes like horizontal localization and inflation, and (ii) differences in the observation error due to the CPSR compression transform. The QOR and CPSR observation errors are different because the compression transform projects the errors onto the leading left singular vectors of the averaging kernel and retains only those components that lie in the range of the averaging kernel.

In summary the steps for obtaining CPSRs from trace-gas retrievals are as follows (assuming the retrieval equation has the same form as Eq. 2):

- Obtain the retrieval and retrieval prior profiles, the averaging kernel, and the observation error covariance for a particular horizontal location.
- Subtract the retrieval prior term $(\mathbf{I} - \mathbf{A})\mathbf{y}_a$ from the retrieval profile \mathbf{y}_r . This yields the QOR as defined by Eq. (3).
- Perform a SVD of the averaging kernel. Form a transform matrix from the left singular vectors associated with the non-zero singular values. Left multiply the QOR by the transpose of the transform matrix. This yields the truncated QOR profile as defined by Eq. (4).
- Left multiply the observation error covariance by the transpose of the transform matrix. Right multiply that matrix product by the transform matrix. This yields the truncated observation error covariance.
- Perform a SVD of the truncated observation error covariance. Scale the left singular vectors with the inverse square root of their respective singular values. Left multiply the truncated QOR profile by the transpose of the scaled left singular vector matrix. This yields the CPSR profile as defined by Eq. (5).
- As a check, left multiply the truncated observation error covariance by the transpose of the scaled left singular vector matrix. Right multiply that matrix product by the scaled left singular vector matrix. The result should be

an identity matrix with rank equal to the number of non-zero CPSRs from the previous step.

- Assimilate the non-zero CPSRs with unitary error variance.

3 WRF-Chem/DART – a regional chemical transport/data assimilation system

The Weather Research Forecasting Model with chemistry/Data Assimilation Research Testbed (WRF-Chem/DART) system is the WRF-Chem chemical transport model (www2.acd.ucar.edu/wrf-chem) coupled with the DART (www.image.ucar.edu/DAReS/DART) ensemble adjustment Kalman filter (Anderson, 2001, 2003) data assimilation system. WRF-Chem/DART is an extension of WRF/DART (www.cawcr.gov.au and references therein). WRF-Chem is the National Center for Atmospheric Research (NCAR) regional Weather Research and Forecasting (WRF) model (www.wrf-model.org) with chemistry. WRF-Chem is a regional model that predicts conventional weather together with the emission, transport, mixing, and chemical transformation of atmospheric trace gases and aerosols. WRF-Chem is collaboratively developed and maintained by the National Oceanic and Atmospheric Administration/Earth System Research Laboratory (NOAA/ESRL), Pacific Northwest National Laboratory (PNNL), and NCAR/Atmospheric Chemistry Observation and Modeling Laboratory (ACOM). WRF-Chem is documented in Grell et al. (2005), discussed in Kukkonen et al. (2012), and has been applied in various research settings (e.g., Pfister et al., 2011, 2013).

DART (Anderson et al., 2009) is a community resource for ensemble data assimilation (DA) research developed and maintained by the NCAR/Data Assimilation Research Section (DAReS). DART is a flexible software environment for studying the interaction between different assimilation methods, observation platforms, and forecast models. WRF-Chem and DART are state-of-the-art tools for studying the impact of assimilating trace gas retrievals on conventional and chemical weather analyses and forecasts.

4 Study period, domain, initial conditions, boundary conditions, emissions, and initial ensemble generation

We conducted continuous cycling experiments with WRF-Chem/DART for the period of 00:00 UTC, 1 June 2008 to 00:00 UTC, 1 July 2008 with 6 h cycling (00:00, 06:00, 12:00, and 18:00 UTC). To facilitate a large number of experiments, we used a reduced ensemble of 20 members and a horizontal resolution of 100 km (101 × 41 grid points). We used 34 vertical levels with a model top at 10 hPa and ~ 15 levels below 500 hPa. WRF-Chem ran with the Model for Ozone and Related Chemical Tracers (MOZART-4) chem-

istry and Goddard Chemistry Aerosol Radiation and Transport (GOCART) model aerosol options (Colarco et al., 2009; Emmons et al., 2010). Ideally for chemical transport forecast experiments we would like an ensemble size of at least 40 members, a horizontal resolution of no larger than 20 km, and a vertical grid with at least 50 levels. We expect our small ensemble/coarse-resolution cycling results, as they pertain to the assimilation of QORs and CPSRs, will apply to larger ensembles with higher resolutions. However, as the vertical resolution increases, the sensitivity to vertical localization may increase (because as the model's vertical resolution increases (i) the vertical solution becomes less smooth and may exhibit greater vertical variability and (ii) the fidelity of vertical localization becomes greater) so that tuning of the vertical localization length may be necessary. For our experiments we used a three-dimensional Gaspari–Cohn type localization with a localization radius half-width of 3000 km in the horizontal and 8 km in the vertical. We conducted sensitivity experiments to determine the appropriate localization settings. Results from the horizontal tests are not discussed. Results from selected vertical localization tests are discussed briefly in Sect. 7.5.

We used NCEP Global Forecast System (GFS) 0.5° six-hour forecasts for the WRF-Chem initial/boundary conditions. Our model domain extends from ~ 176 to ~ 50° W and from ~ 7 to ~ 54° N. We used the WRF preprocessing system (WPS) to interpolate the GFS forecasts to our domain and generate the deterministic boundary conditions. We used the WRF data assimilation system (WRFDA) (http://www2.mmm.ucar.edu/wrf/users/wrfda/Docs/user_guide_V3.7/WRFDA_Users_Guide.pdf) to generate the initial meteorology ensemble.

For the chemistry initial and lateral boundary conditions, we used global simulations from the NCAR MOZART-4 model. The fire emissions came from the Fire Inventory from NCAR (FINNv1; Wiedinmyer et al., 2011), and the Model of Emissions of Gases and Aerosols from Nature (MEGAN; Guenther et al., 2012) calculated the biogenic emissions as part of the WRF-Chem forecast. The anthropogenic emissions were based on the US Environmental Protection Agency's (EPA's) 2005 National Emissions Inventory (NEI-2005). We used or adapted existing ACOM/WRF-Chem utilities (<https://www2.acom.ucar.edu/wrf-chem/wrf-chem-tools-community>) to generate the initial chemistry ensembles with a Gaussian distribution from a specified mean and standard deviation. That distribution was truncated at the tails to include 95 % of the distribution. Similar utilities were used to generate the emission ensembles. We excluded the distribution tails to avoid the potential for the extreme values to cause numerical problems in the chemistry algorithms. Although we recognize that the assimilation cycling results may be sensitive to the emission perturbation horizontal correlation lengths (e.g. Pagowski and Grell, 2012), this was not particularly relevant to our study

so we set the horizontal and vertical correlation lengths to zero.

5 Meteorology observations and satellite trace gas retrievals

At each cycle time, depending on the experiment we assimilated meteorology and/or chemistry observations with the DART ensemble adjustment Kalman filter (EAKF) and then advanced the analysis ensemble to the next cycle time with WRF-Chem. The 6 h forecast ensemble was then used as the first guess for the next ensemble DA step.

We assimilated conventional meteorological observations and CO trace gas retrievals from MOPITT. The meteorological observations were NCEP automated data processing (ADP) upper air and surface observations (PREPBUFR observations). They included air temperature, sea level pressure, surface winds, dew point temperature, sea surface temperature, and upper level winds from various observing platforms. We refer to those observations as the *MET OBS*.

We also assimilated MOPITT partial column/profile CO retrievals. MOPITT is an instrument flying on NASA's Earth Observing System Terra spacecraft. MOPITT's spatial resolution is 22 km at nadir, and it sees the earth in 640 km wide swaths. MOPITT uses gas correlation spectroscopy to measure CO in a thermal-infrared (TIR) band near 4.7 μm and a near-infrared (NIR) band near 2.3 μm . TIR radiances are sensitive to CO in the middle and upper troposphere while NIR measures the CO total column. Worden et al. (2010), Deeter (2011), and Deeter et al. (2012, 2013) showed that the sensitivity to CO in the lower troposphere is significantly greater for retrievals exploiting simultaneous TIR and NIR than for retrievals based on TIR alone. MOPITT started data collection in March 2000. We used the MOPITT v5 TIR/NIR products described in Deeter et al. (2013). We refer to the MOPITT observations as the *CHEM OBS*.

The retrieval error covariance \mathbf{E}_r associated with each MOPITT CO retrieval profile is provided as part of the data product. That error covariance is derived by the retrieval process based on a specified a priori error covariance \mathbf{E}_a . Under the optimal estimation theory of Rodgers (2000) \mathbf{E}_r is related to \mathbf{E}_a through the averaging kernel \mathbf{A} by $\mathbf{E}_r = (\mathbf{I} - \mathbf{A})\mathbf{E}_a$. The measurement error in retrieval space \mathbf{E}_m is also related to \mathbf{E}_a and \mathbf{A} by $\mathbf{E}_m = (\mathbf{I} - \mathbf{A})\mathbf{E}_a\mathbf{A}^T$. Generally for retrieval data sets, \mathbf{E}_a , \mathbf{E}_r , and \mathbf{E}_m are non-diagonal.

6 Experimental design

We conducted two basic experiments: (i) a control experiment where we assimilated only MET OBS (MET DA); and (ii) a chemical data assimilation experiment where we assimilated MET OBS and MOPITT CO partial column retrievals in the form of QORs (MOP QOR). In addition we conducted an experiment where we converted the CHEM OBS to CP-

SRs and assimilated the CPSRs (MOP CPSR). We also conducted sensitivity experiments where we (i) zeroed the observation error covariance cross-correlations (MOP NROT) – as opposed to using a SVD transformation for diagonalization, and (ii) applied vertical localization (MOP LOC). The suite of experiments is summarized in Table 1.

For all experiments we used (i) DART horizontal and vertical localization – Gaspari–Cohn localization with a localization radius half-width of 3000 km in the horizontal and 8 km in the vertical, (ii) DART prior adaptive inflation, (iii) no posterior inflation, (iv) full interaction between all observations and all state variables – i.e., MET OBS update chemistry state variables and CHEM OBS update meteorology state variables (joint assimilation of MET and CHEM OBS), (v) DART clamping (i.e., the imposition of a minimum threshold) on chemistry state variables to constrain the posterior ensemble members to be positive, and (vi) the reported MOPITT retrieval error covariance as the observation error covariance to account for unrepresented error sources such as representativeness error.

For the MOP QOR experiment, the MOPITT CO retrievals were converted to QORs using an algorithm similar to that described by Migliorini et al. (2008) except we did not perform their second forecast error covariance-based filtering.

7 Results

7.1 The control and chemical data assimilation experiments

The MET DA and MOP QOR experiments are intended to identify the impact of assimilating chemistry observations. Figure 1a shows shaded contours of CO in parts per billion (ppb) at ~ 1000 hPa from the 6 h forecast valid at 00:00 UTC, 29 June 2008 and compares the MET DA and MOP QOR experiments. It shows that over the course of MOP QORs the assimilation of CO retrievals reduced the (i) positive CO bias found in polluted areas of MET DA (i.e., metropolitan areas with high-CO emissions – San Francisco, Los Angeles, Chicago, and the northeast USA), and (ii) negative CO bias found in nonpolluted areas in MET DA (Hawaii, east Pacific, southeast USA, and Baja). The MET DA biases could result from model errors such as (i) emission errors – CO emissions too high in polluted areas and too low in nonpolluted areas, (ii) transport errors – insufficient CO transport away from polluted areas and insufficient transport toward nonpolluted areas, and/or (iii) chemistry errors – CO destruction too weak in polluted areas and too strong in nonpolluted areas. The MET DA biases could also result from initial/boundary condition errors that were corrected by the assimilation of MOPITT CO in MOP QORs.

Figure 1b shows the assimilated CO retrievals for the 18:00 UTC, 28 June 2008 update cycle in the upper panel and the corresponding increments in the lower panel. Compari-

Table 1. Summary of the WRF-Chem/DART Forecast/Data Assimilation Experiments.

Experiment	Assimilate meteorology observations	Assimilate MOPITT CO QORs	Assimilate MOPITT CO CPSRs	Use error covariance zeroing	Use vertical localization
MET DA	Yes	No	No	No	No
MOP QOR	Yes	Yes	No	No	No
MOP CPSR	Yes	No	Yes	No	No
MOP NROT	Yes	Yes	No	Yes	No
MOP LOC	Yes	Yes	No	No	Yes

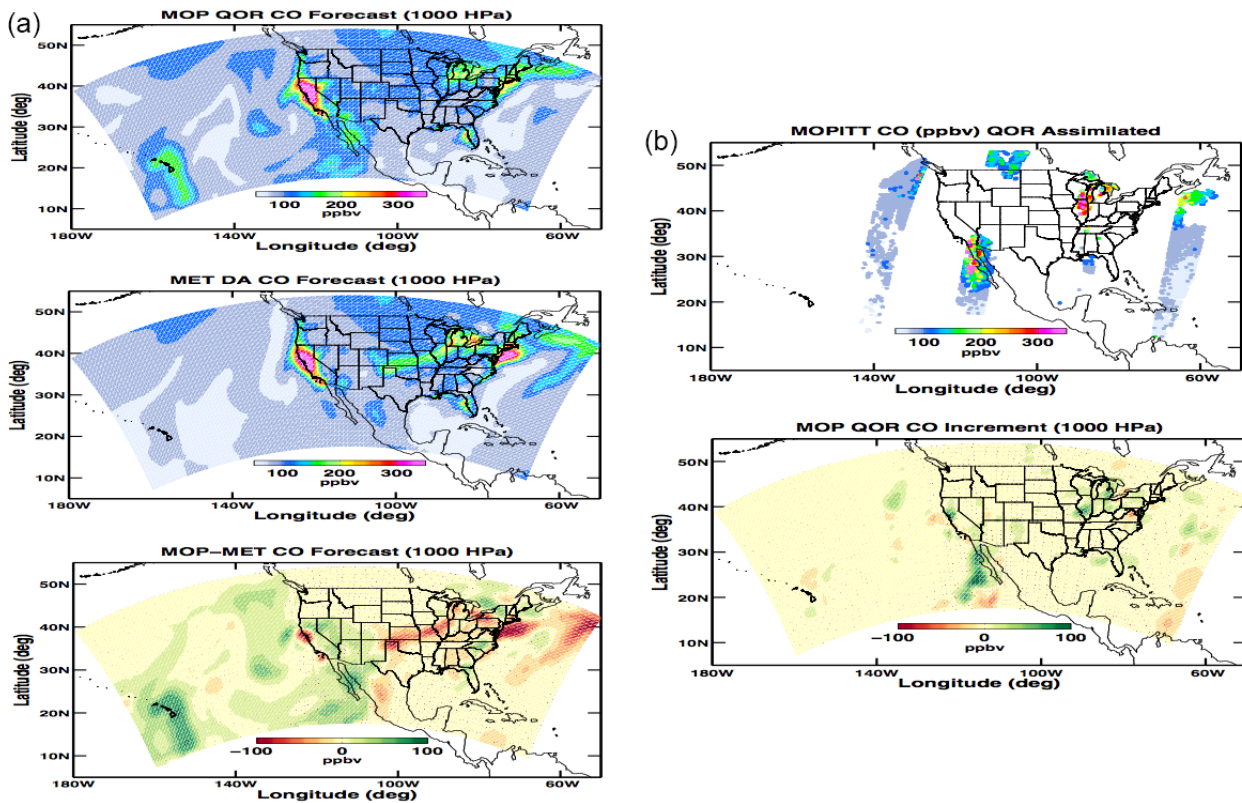


Figure 1. (a) Shaded contours of CO in ppb for the MOP QOR (upper panel) and MET DA (middle panel) experiments for the first model level above the surface (~ 1000 hPa) from the 6 h forecast valid on 18:00 UTC, 28 June 2008. The lower panel shows the difference contours for those experiments (MOP QOR – MET DA). The shaded area represents the WRF-Chem domain. (b) The upper panel shows the assimilated MOPITT CO retrievals between the surface and 900 hPa for 18:00 UTC, 28 June 2008. The lower panel shows the associated assimilation increment.

son of those panels shows that the assimilation step adjusted the CO concentrations primarily along the satellite observation paths, which is a consequence of assimilating sparse observations. The DA adjustments in Fig. 1b are generally consistent with the differences between MOP QORs and MET DA in Fig. 1a (CO increases in nonpolluted areas – east of San Francisco, the southeast USA, and Baja). However, that is a general statement because the MOP QOR – MET DA differences are partially related to the impact of assimilating CO observations during the preceding assimilation cycle and partially related to the impact of assimilating all the CO

observations since the beginning of the cycling experiment (~ 100 cycles). Consequently, there are locations where the signs of the MOP QOR – MET DA differences are different from the signs of the increments (e.g. southwest of lakes Michigan and Huron and over the Ohio River valley and San Francisco Bay). The sense of those sign differences is not an indication of relative forecast accuracy but that the (i) impact from assimilating CO during the preceding cycle was similar to that from assimilating CO throughout the cycling experiment (same signs), and (ii) impact from assimilating CO during the preceding cycle was different to that from as-

simulating CO throughout the cycling experiment (different signs).

Figure 2 shows time series of the domain average CO from the MET DA and MOP QOR experiments in retrieval and state space. The dots represent the retrieval space results where the cool colors (blue and black) show the forecasts, and the warm colors (red and magenta) show the analyses. The green dots represent the MOPITT retrievals. The solid lines show state-space results. Figure 2 has several interesting results. First, MET DA had a negative bias of ~ 10 ppb in retrieval space. Second, assimilation of MOPITT CO reduced that bias by ~ 5 ppb. Finally, in state space MOP QORs increased the mean CO by ~ 5 ppb. As discussed below, those results are consistent with Fig. 1, which shows a large number of nonpolluted areas in MET DA with a negative bias and a small number of polluted areas with a positive bias.

Figure 3 shows vertical profiles of the time (00:00 UTC, 25 June 2008 to 00:00 UTC, 29 June 2008) and horizontal domain average CO in retrieval space. It shows that the MOPITT profile had greater vertical variability (moderate CO near the surface, low CO in the middle troposphere: 500–400 hPa, high CO in the upper troposphere: 300–200 hPa, and low CO near the tropopause: 200–100 hPa) than the MET DA and MOP QOR profiles. It also shows that the assimilation of MOPITT CO had positive impacts throughout the troposphere with the greatest improvement in the upper troposphere. Figure 3 shows that there were differences in the MOP QOR/MET DA bias reduction between: (i) the upper and lower troposphere (greater magnitude negative bias reduction in the upper troposphere and lesser magnitude positive bias reduction in the lower troposphere), and (ii) the forecast and the analysis (greater bias reduction in the analysis than in the forecast). Those results expand our understanding of the bias in Figs. 1 and 2. In Fig. 3 the forecast and analysis show greater bias reduction in the upper troposphere. That suggests that the domain averages in Fig. 2 were dominated by bias reductions in the upper troposphere. Figure 3 also suggests that bias reductions in the lower troposphere were dominated by the reduction of the positive bias in the polluted areas of Fig. 1. Those results suggest that the following model errors (as opposed to initial/boundary condition errors) caused the biases: (i) the near-surface biases were likely caused by the CO emissions being too high in polluted areas and too low in nonpolluted areas, (ii) the positive biases in the lower middle troposphere (~ 600 hPa) were likely caused by erroneously large vertical CO fluxes from the near surface to the lower middle troposphere and/or too little CO destruction, and (iii) the negative biases in the upper troposphere were likely caused by erroneously small vertical CO fluxes and/or too much CO destruction. We reach the conclusion regarding model error versus initial/boundary condition (IC/BC) error because Fig. 3 shows that the bias reduction in the lower troposphere is greater for the analyses than for the forecasts. That suggests that following the assimilation of

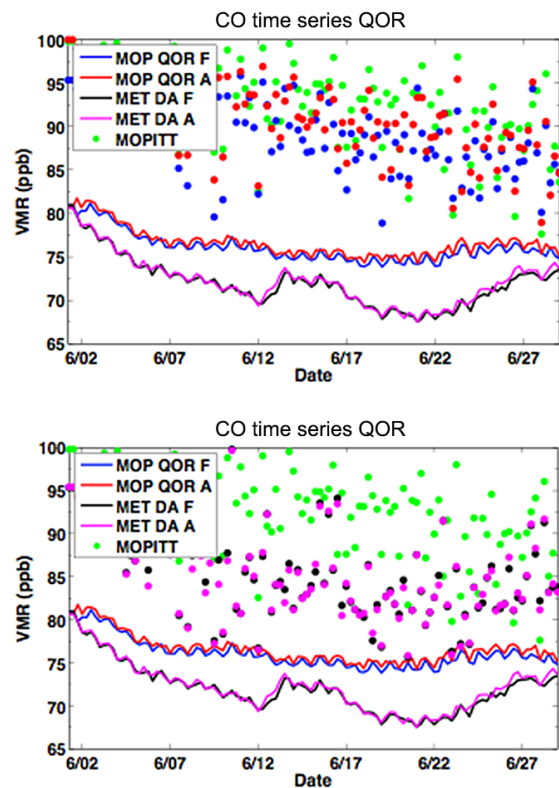


Figure 2. Time series of the domain average CO from the MOP QOR and MET DA experiments. The red and magenta dots show the domain average CO in retrieval space for the MOP QOR and MET DA analyses denoted in the legend by “A”. The blue and black dots show the domain average CO in retrieval space for the MOP QOR and MET DA forecasts denoted in the legend by “F”. The green dots show the domain average MOPITT CO retrievals. They are the same in both panels and are included for reference. The solid lines show the domain average CO in model space with the same color scheme as used for the analyses and forecasts in retrieval space. The solid lines are the same in both panels are also included for reference.

MOPITT CO in MOP QOR, the CO IC/BCs have improved relative to MET DA. Then during the course of model integration the bias increases. Thus, we conclude that model error is a more likely cause of the bias.

Lastly, we tested the null hypothesis that the difference between the MET DA and MOP QOR time series results was zero ($H_0: \text{MOP QOR} - \text{MET DA} = 0$) against an alternative hypothesis that the difference was not zero ($H_A: \text{MOP QOR} - \text{MET DA} \neq 0$). We used the retrieval-space time series from Fig. 2 and the large sample parametric test for the difference between two means from a normal distribution. The test statistic was $Z = \frac{\bar{Y}_1 - \bar{Y}_2}{\sqrt{\sigma_1^2/n_1 - \sigma_2^2/n_2}}$ where \bar{Y}_1 , σ_1^2 , and n_1 denote the sample mean, sample variance, and number of samples for the MOP QOR experiment, respectively; \bar{Y}_2 , σ_2^2 , and n_2 denote the analogous sample statistics for the MET DA experiment; and $n_1 = n_2 = 104$. The rejection criteria

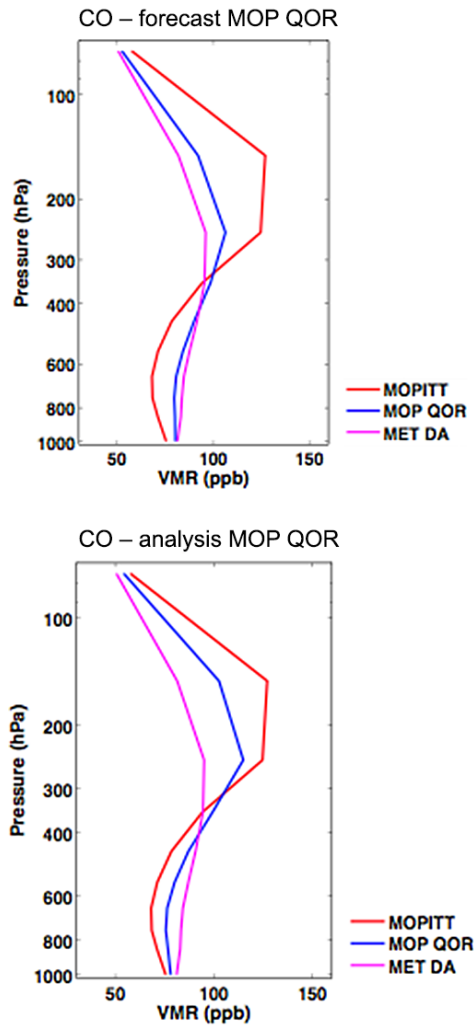


Figure 3. Vertical profiles of time/horizontal domain average CO from the MOP QOR and MET DA experiments for 00:00 UTC, 25–29 June 2008. The results are in MOPITT retrieval space. The red profiles represent the MOPITT retrievals. Otherwise the color of the lines corresponds to the legend. *forecast* is the assimilation prior, and *analysis* is the assimilation posterior.

was “Cannot handle” as spaceZ cannot handle” as space $> z_{\alpha/2}$, where $\alpha = 0.05$ and $z_{\alpha/2} = 1.96$ for a two-tailed test at the 95 % confidence level. We were able to reject the null hypothesis. Based on that result, we conclude that assimilation of MOPITT CO retrievals significantly changed the WRF-Chem/DART CO forecasts and analyses. When measured against MOPITT, those changes were a significant improvement.

7.2 Assimilation of compact phase space retrievals

Next we study the assimilation of CPSRs as described in Sect. 2 but first review some CPSR attributes. Figure 4a shows vertical profiles of CPSR characteristics averaged for

the MOPITT retrieval domain at 18:00 UTC, 28 June 2008. The blue curves represent the MOPITT CO retrievals (MOP-Rets). Those curves have reduced vertical structure due to the units ($\log_{10}(\text{VMR})$ as opposed to VMR). After conversion from $\log_{10}(\text{VMR})$ to VMR, MOP-Rets has greater vertical structure and resembles the MOPITT profiles in Fig. 3. The black curves represent the MOPITT CO QORs (MOP-QOR) as defined by Eq. (3). MOP-QORs differs from MOP-Rets in that they have maxima near the surface and upper troposphere and a minimum in the middle troposphere. The green curves represent the truncated profiles, which are obtained by (i) projecting the full retrieval profile or the QOR profile onto the leading left singular vectors of the associated averaging kernel to get the projection coefficients (e.g., $\mathbf{c}_r = \mathbf{U}_0^T \mathbf{y}_r$, where \mathbf{c}_r is the projection coefficient vector for the full retrieval and $\mathbf{c}_{\text{qor}} = \mathbf{U}_0^T (\mathbf{y}_r - (\mathbf{I} - \mathbf{A}) \mathbf{y}_a - \boldsymbol{\varepsilon}) = \mathbf{U}_0^T \mathbf{y}_{\text{qor}}$, where \mathbf{c}_{qor} is the coefficient vector for the QOR profile – see Eq. 4 in Sect. 2), and (ii) performing the inverse projection by multiplying the leading singular vectors by their respective projection coefficients and summing those dot products (e.g., $\hat{\mathbf{y}}_r = \mathbf{U}_0 \mathbf{c}_r$ is the truncated retrieval profile – denoted MOP-Trc and $\hat{\mathbf{y}}_{\text{qor}} = \mathbf{U}_0 \mathbf{c}_{\text{qor}}$ is the truncated QOR profile – denoted QOR-Trc). The forward transform in (i) is analogous to the first part of the CPSR transform in Eq. (4). The inverse transform in (ii) brings the result of forward transform in (i) back to state space. The inverse transform is not part of the CPSR algorithm.

In Fig. 4a the residuals are defined as the difference between the full and truncated profiles (e.g., $\mathbf{y}_r - \hat{\mathbf{y}}_r$ is the full retrieval residual – denoted MOP-Res and $\mathbf{y}_{\text{qor}} - \hat{\mathbf{y}}_{\text{qor}}$ is the QOR residual – denoted QOR-Res). If the full profiles project completely onto the leading singular vectors, the residuals are zero. The upper panel of Fig. 4a shows that the transform in (i) has the greatest impact near the surface and the upper troposphere and the least impact in the middle troposphere. When the truncation residuals are nonzero, the original profiles contain components that are not in the range of the averaging kernel. That always indicates a contribution from the retrieval prior term $(\mathbf{A} - \mathbf{I}) \mathbf{y}_a$. However, a zero residual does not always indicate that the contribution from the retrieval prior term has been removed. For QOR residuals, the retrieval prior term contribution is completely removed. For the retrieval residuals, the retrieval prior term contribution may not be completely removed. When components of the retrieval prior term lie in the range of the averaging kernel, they cannot be removed by the transform in (i) and are therefore not included in the residual. For example in the upper panel of Fig. 4a the similarity between MOP-Trc and QOR-Ret shows that the MOPITT retrieval was strongly influenced by the retrieval prior and that most of the prior contribution was removed by the transform in (i). That also shows that most of the prior contribution was not in the range of the averaging kernel. However, not all was outside the range, and the difference between MOP-Trc and QOR-Ret

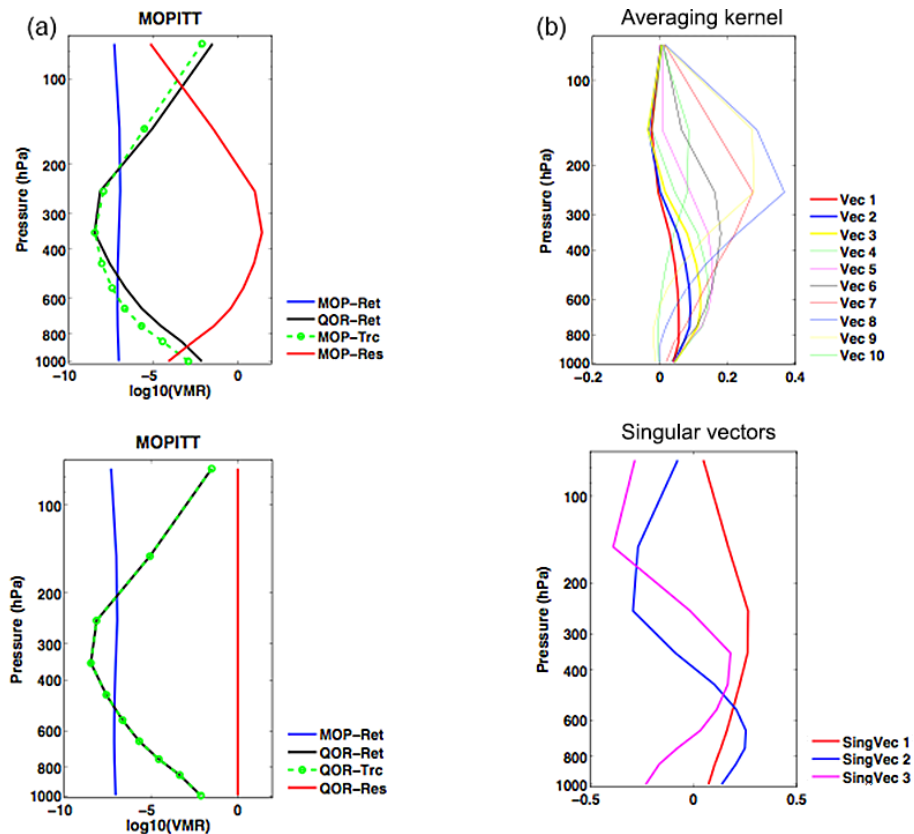


Figure 4. (a) Horizontal domain average of the full and truncated terms in the retrieval equation for 18:00 UTC, 28 June 2008. MOP-Ret, MOP-Trc, and MOP-Res are the MOPITT retrieval, truncated retrieval, and residual profiles, respectively. QOR-Ret is the MOPITT QOR profile, QOR-Trc is the truncated MOPITT QOR profile, and QOR-Res is the MOPITT QOR residual profile. (b) Horizontal domain average of the MOPITT averaging kernel profiles in the upper panel and leading left singular vectors of those averaging kernels in the lower panel for 18:00 UTC, 28 June 2008.

shows that most was inside the range. This analysis shows that the influence of the retrieval prior term cannot be completely removed by projecting the retrieval onto the range of the averaging kernel. The results show that it is necessary to use the Migliorini et al. (2008) quasi-optimal subtraction in Eq. (2) to remove the retrieval prior contribution. Comparison of QOR-Ret and QOR-Trc in the lower panel of Fig. 4a shows that QOR-Ret lies completely within the range of the averaging kernel. That result was expected from the discussion of Eqs. (3) and (4).

In summary Fig. 4a shows the state space impacts from applying the Migliorini et al. (2008) quasi-optimal subtraction and the CPSR transform in (i). It also shows that the quasi-optimal subtraction was necessary to remove the influence of the retrieval prior. Thus, in CPSRs the quasi-optimal subtraction removes the influence of the retrieval prior, and projection onto the leading singular vectors of the averaging kernel provides the data compression.

In Fig. 4a the average number of leading singular vectors was ~ 2.3 . CPSRs reduced the number of observations by ~ 7.7 per MOPITT profile. After thinning there were

$\sim 30\,000$ MOPITT profiles per assimilation cycle. That implies a CPSR reduction of $\sim 281\,000$ retrievals or $\sim 80\%$ per cycle. On application the actual reduction was less because the number of non-retrieval observations was not reduced. As an example when assimilating MET OBS and CHEM OBS we found a reduction of $\sim 35\%$ in the computation cost. That is a wall clock time reduction based on NCAR's 1.5-petaflop high-performance IBM Yellowstone computer with 32 tasks, and 8 tasks per node. We expect similar reductions for other computing configurations.

In Fig. 4b we examine the vertical structure of the CPSRs. The upper panel shows the retrieval domain average of the MOPITT averaging kernel profiles from 18:00 UTC, 28 June 2008. The lower panel shows the domain average of the leading left singular vectors from SVDs of the averaging kernel profiles in the upper panel. Comparison of the upper and lower panels shows that while the singular vector and averaging kernel profiles are similar it is not possible to associate a specific singular vector profile with a specific averaging kernel profile or with a group of profiles. However, the averaging kernel of singular vectors show the sensitiv-

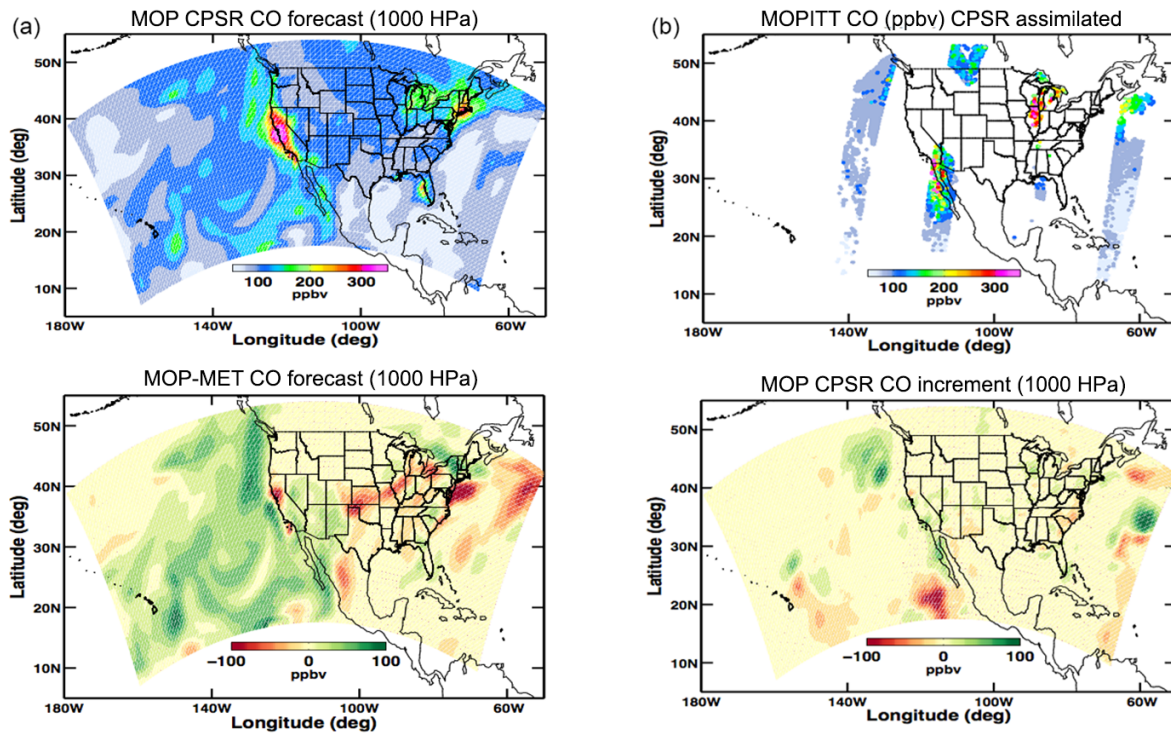


Figure 5. (a) Same as Fig. 1a except for the MOP CPSR experiment and the middle panel from Fig. 1a, the MET DA experiment is not plotted. (b) Same as Fig. 1b except for the MOP CPSR experiment.

ity of the associated CPSR to the true CO profile. The first singular vector shows positive sensitivity to the entire CO profile with greater sensitivity in the lower and middle troposphere and greatest sensitivity in the upper middle troposphere. The second singular vector shows positive sensitivity in the lower troposphere and negative sensitivity in the upper troposphere. Lastly, the third singular vector resembles the first singular vector with greatest positive sensitivity in the upper middle troposphere but with negative sensitivity in the lower and upper troposphere. Those characteristics are consistent with the MOPITT TIR and NIR joint sensitivities documented by Worden et al. (2010), Deeter (2011), and Deeter et al. (2012, 2013). It should be noted that the sign of the singular vectors in the Fig. 4b is arbitrary because the left and right singular vectors can be jointly multiplied by negative one and still qualify as singular vectors. However, when multiplied by one sign the singular vector may have physical meaning, and when multiplied by the other it may not. For our application, the sign that made the vertical structure of the singular vectors most similar to that of the averaging kernel had physical meaning. Therefore, in Fig. 4b we chose the sign that made the singular vector profile most consistent with the averaging kernel profiles.

To test the benefit of assimilating CPSRs we converted the MOPITT CO retrievals to CPSRs and repeated the MOP QOR experiment (called MOP CPSR). Those results are shown in Figs. 5–7. Conceptually the MOP CPSR results

should be similar to the MOP QOR results in Figs. 1–3. Practically, the results are different due to (i) the effect of DA sub-processes like horizontal localization and inflation, and (ii) differences in the observation error caused by the CPSR compression transform. Comparison of the contour maps in Figs. 1a and 5a shows that MOP CPSR provided similar adjustments to MOP QOR but they were of greater magnitude and larger area (the MET DA result was not plotted in Fig. 5a because it would be the same as in Fig. 1a). The general trend from Fig. 1a that the assimilation of CO retrievals reduced the positive CO bias in polluted areas and the negative bias in nonpolluted areas appears in Fig. 5a. Comparison of Figs. 1b and 5b shows that MOP CPSR generally assimilated the same CO retrievals as MOP QOR, but the CPSR increments were of greater magnitude and more widely dispersed. Comparison of the time series plots in Figs. 2 and 6 shows that there were slightly greater bias reductions for MOP CPSR than MOP QOR. MOP CPSR reduced the CO negative bias in retrieval space by ~ 8 ppbv and increased the mean CO in state space by ~ 10 ppbv. Those improvements are also seen from a comparison of the vertical profiles in Figs. 3 and 7, which shows that MOP CPSR produced greater bias reductions for the forecast and analysis throughout the troposphere. As in MOP QOR the MOP CPSR improvements were greater in the upper troposphere than in the lower troposphere. The MOP CPSR results from Fig. 7 provide further support for our suggestion that the domain average bias re-

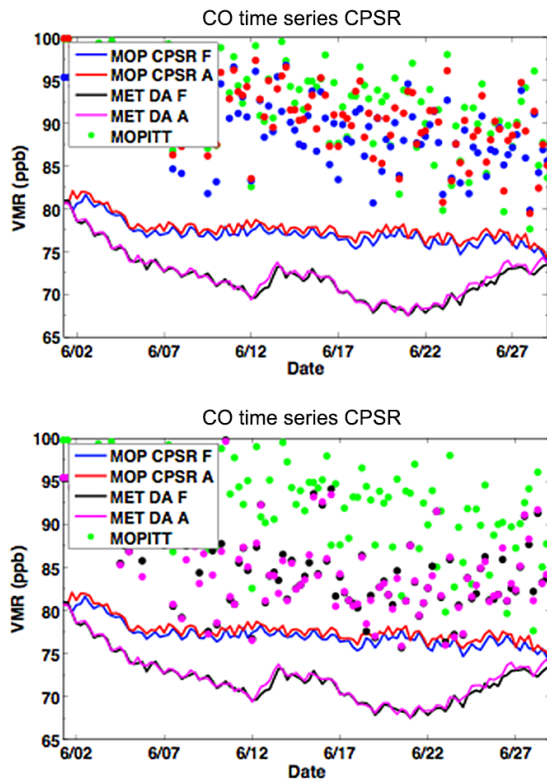


Figure 6. Same as Fig. 2 except for the MOP CPSR experiment.

ductions in Figs. 2 and 6 were due to bias reductions in the upper troposphere because the greater bias reductions in the upper troposphere of Fig. 7 (compared to Fig. 3) provided greater bias reductions in Fig. 6 (compared to Fig. 2). In summary Figs. 5–7 confirm our analysis of Figs. 1–3 and show that assimilation of CPSRs produced results that were similar to or better than those from the assimilation of QORs at two-thirds the computational cost. We also conducted significance testing for MOP CPSR similar to that for MOP QOR and were able to reject the null hypothesis that there was no difference between the MOP CPSR and MET DA time series in Fig. 6.

7.3 Verification against MOPITT retrievals

We calculated verification statistics (bias, root mean square error (RMSE), and coefficient of determination (R^2)) for the 6 h forecasts from all experiments based on the time series results in Figs. 2 and 6. Those statistics are plotted in Fig. 8. Generally, Fig. 8 shows that the assimilation of MOPITT CO improved model performance for all metrics when compared against the MOPITT retrievals. Figure 8 also shows that RMSE was dominated by the bias and that the differences in the statistics for the different treatments (assimilating QORs, CPSRs, cross-covariance zeroing, and vertical localization) were generally negligible except for cross-covariance zero-




Figure 7. Same as Fig. 3 except for the MOP CPSR experiment.

ing. We now discuss the cross-covariance zeroing and vertical localization experiments.

7.4 Observation error covariance diagonalization through zeroing of the cross-correlations

One method used to diagonalize the observation error covariance is zeroing of the cross-correlations (see the Introduction to Migliorini et al., 2008). The uncertainty of the error covariance and the practice of adjusting the observation error variance to tune ensemble DA strategies are used to justify the zeroing. As noted by Anderson (2001) and applied by Migliorini et al. (2008), a more aesthetic and mathematically correct approach is to apply a variance maximizing rotation based on a SVD of the error covariance. In this section we compare those two error covariance diagonalization methods. Recall that MOP QOR used an SVD-based rotation to diagonalize the error covariance. We conducted a companion experiment MOP NROT where we used cross-correlation zeroing. The assimilation/forecast plots are not shown because




Figure 8. Verification statistics for all experiments in MOPITT retrieval space. The blue curve is the bias (model – observation), the red curve is the root mean square error (RMSE), and the magenta curve is the coefficient of determination. The experiments are described in the text and summarized in Table 1.

it is not a central theme of this paper. However, we include the verification statistics in Fig. 8. Significance testing and scores from assimilation of MOPITT CO show that SVD-based diagonalization produced significantly greater forecast skill compared to cross-correlation zeroing. Based on that result we conclude that the second SVD-based rotation is a necessary step in our definition of CPSRs.

7.5 Vertical localization and phase space retrievals

Ensemble data assimilation generally uses localization to remove spurious correlations that may occur from under-sampling. Localization limits the horizontal and vertical spatial scales on which the observations impact the posterior. Vertical localization may be inappropriate when assimilating phase space retrievals because $\sim 80\%$ of the vertical variation in the retrieval is described by the first leading singular vector of the averaging kernel (the basis function for the phase space transform) and that vector is nearly independent of height (see Fig. 4b). Nevertheless, if vertical localization is appropriate then the question becomes *how to do it* because phase space retrievals are not associated with a unique vertical location. One solution assumes that phase space retrievals are associated with the level of maximum sensitivity in the transformed averaging kernel, i.e., the averaging kernel after applying the compression and diagonalization transforms

discussed in Sect. 2. We applied such localization to MOP QOR (in the MOP LOC experiment) and found that results from the two experiments were similar. Therefore, we do not present the assimilation/forecast plots but include the verification statistics in Fig. 8. Comparison of the verification scores in Fig. 8 for MOP LOC with those from the other experiments (MOP QOR and MOP CPSR) shows that vertical localization did not substantially alter the results. We experimented with different vertical localization lengths and found similar results. We are unsure whether this is a general result and are continuing to investigate vertical localization.

8 Summary and conclusions

In this paper we incorporated WRF-Chem into DART and assimilated MOPITT CO trace gas retrievals. We also introduced the assimilation of compact phase space retrievals (CPSRs). CPSRs are preprocessed trace gas retrievals that have (i) the influence of the retrieval prior removed, (ii) data compression, (iii) SVD-based error covariance diagonalization, and (iv) unit error variance scaling. We showed that assimilation of CPSRs is an efficient alternative to assimilation of quasi-optimal retrievals (QORs) that provided substantial reductions in computation time ($\sim 35\%$) without degrading the analysis fit or forecast skill.

We presented results from month-long (00:00 UTC, 1 June 2008 to 18:00 UTC, 31 June 2008) cycling experiments where we assimilated conventional meteorology and MOPITT CO retrievals. For MOP QOR the time series plots in Fig. 2 showed that MET DA had a negative bias of ~ 10 ppb. The assimilation of MOPITT CO in MOP QOR reduced that bias by ~ 5 ppb. The vertical profile plots in Fig. 3 showed that assimilation of MOPITT CO improved the CO analysis fit and forecast skill throughout the troposphere when compared to MET DA. We also used traditional skill metrics (bias, RMSE, and R^2) to quantify the impact of assimilating CO retrievals. Those results showed that bias dominated the RMSE and that assimilation of CO retrievals improved WRF-Chem performance. Specifically, MOP QOR significantly improved the WRF-Chem CO forecast skill for all three metrics.

Next we focused on making the assimilation of retrievals computationally efficient and introduced compact phase space retrievals. CPSRs advance the work of Joiner and DaSilva (1998) and Migliorini et al. (2008) by describing an easily applied methodology to achieve data compression for phase space retrievals. Conceptually, the assimilation of QORs and CPSRs should yield similar results except for the effects of (i) assimilation sub-processes like localization and inflation and (ii) different observation errors due to the CPSR compression transform. Nevertheless, our CPSR approach is different from that of Migliorini et al. (2008): (i) we perform two transforms – a compression transform and a diagonalization transform, they perform two diagonalization transforms;

(ii) we identify and assimilate all linearly independent information observed by the instrument, they may discard linearly independent information – some because their transform vectors are not necessarily a basis for the space of QORs and some because their uncertainty analysis discards some information that lies in the range of their transformed averaging kernel; (iii) our diagonalization transform is analogous to their first diagonalization transform except we diagonalize the compressed observation error covariance and they diagonalize the untransformed observation error covariance; and (iv) we rely on the assimilation system to decide how much weight to give the transformed observations and require no information from the forecast ensemble, and they use the forecast ensemble to decide which observations to discard.

MOP CPSR maps in Fig. 5 showed that assimilation of CPSRs placed CO hot spots in the same locations as MOP QOR but they were of greater magnitude and larger area. The time series and vertical profile plots showed that those differences generally represented analysis and forecast improvements. Skill metrics for MOP CPSR showed that when compared to MOP QOR, the assimilation of CPSRs slightly improved the forecast skill for all metrics, and when compared to MET DA it significantly improved the forecast skill for all metrics. Based on those results we conclude that the assimilation of CPSRs performed as well or better than the assimilation of QORs at a substantially reduced computational cost ($\sim 35\%$ reduction in computation time).

Collectively our analysis of the MOP QOR and CPSR results in Figs. 1–3 and 5–7 suggested that (i) in the lower troposphere MET DA had a negative CO bias in polluted areas and a positive bias in nonpolluted areas (Figs. 1 and 5) and (ii) bias reductions in the domain average retrieval space CO were due to reductions in the negative CO bias in the upper troposphere (Figs. 2, 3, 6, and 7). We proposed three causes for the CO biases: (i) emission errors – overestimation of CO emissions in polluted areas and underestimation in non-polluted areas, (ii) transport errors – too much CO transport from the near surface to the lower troposphere and too little transport from the lower to upper troposphere, and (iii) chemistry errors – too little CO destruction in the near surface and lower troposphere and too much destruction in the upper troposphere.

We expect that CPSRs have the potential for broad operational application. CPSRs can be easily obtained from retrievals derived from any optimal estimation algorithm. They can be used to assimilate retrievals with correlated or uncorrelated errors for any sequential assimilation methodology (both Kalman filter and variational-based algorithms). Due to their ease of derivation, flexibility, and potential for large reductions in assimilation computation time, CPSRs should facilitate the efficient assimilation of dense geostationary observations.

Acknowledgements. NCAR is sponsored by the National Science Foundation (NSF). Any opinions, findings and conclusions or recommendations expressed in this publication are those of the authors and do not necessarily reflect the views of NSF. This research was also sponsored by NASA grants NNX11A110G and NNX10AH45G. We gratefully acknowledge the anonymous reviewers, Chris Snyder, and Louisa Emmons for their thorough reviews of this manuscript and for providing many constructive comments. We also acknowledge Helen Worden for assistance with the MOPITT data, and Jerome Barre for assistance with coding the WRF-Chem initial and boundary condition perturbations.

Edited by: A. Lauer

References

- Anderson, J. L.: An ensemble adjustment Kalman filter for data assimilation, *Mon. Weather Rev.*, 129, 2884–2903, 2001.
- Anderson, J. L.: A local least squares framework for ensemble filtering, *Mon. Weather Rev.*, 131, 634–642, 2003.
- Anderson, J. L., Hoar, T., Raeder, K., Liu, H., Collins, N., Torn, R., and Arellano, A.: The Data Assimilation Research Testbed: A community facility, *B. Am. Meteorol. Soc.*, 90, 1283–1296, 2009.
- Bei, N., de Foy, B., Lei, W., Zavala, M., and Molina, L. T.: Using 3DVAR data assimilation system to improve ozone simulations in the Mexico City basin, *Atmos. Chem. Phys.*, 8, 7353–7366, doi:10.5194/acp-8-7353-2008, 2008.
- Colarco, P., da Silva, A., Chin, M., and Diehl, T.: Online simulation of global aerosol distributions in the NASA GEOS-4 model and comparisons to satellite and ground-based aerosol optical depth, *J. Geophys. Res.*, 115, D14207, doi:10.1029/2009JD012820, 2010.
- Deeter, M. N.: MOPITT (Measurement of Pollution in the Troposphere) version 5 Product User's Guide, available at: www.aom.ucar.edu/mopitt/v5_users_guide_beta.pdf (last access: 1 March 2016), 2011.
- Deeter, M. N., Worden, H. M., Edwards, D. P., Gille, J. C., and Andrews, A. E.: Evaluation of MOPITT retrievals of lower-tropospheric carbon monoxide over the United States, *J. Geophys. Res.*, 117, D13306, doi:10.1029/2012JD017553, 2012.
- Deeter, M. N., Martinez, A. S., Edwards, D. P., Emmons, L. K., Gille, J. C., Worden, H. M., Pittman, J. V., Daube, B. C., and Wofsy, S. C.: Validation of MOPITT version 5 thermal-infrared, near-infrared, and multispectral carbon monoxide profile retrievals for 2000–2011, *J. Geophys. Res.-Atmos.*, 118, 6710–6725, 2013.
- Emmons, L. K., Walters, S., Hess, P. G., Lamarque, J.-F., Pfister, G. G., Fillmore, D., Granier, C., Guenther, A., Kinnison, D., Laepple, T., Orlando, J., Tie, X., Tyndall, G., Wiedinmyer, C., Baughcum, S. L., and Kloster, S.: Description and evaluation of the Model for Ozone and Related chemical Tracers, version 4 (MOZART-4), *Geosci. Model Dev.*, 3, 43–67, doi:10.5194/gmd-3-43-2010, 2010.
- Gaubert, B., Coman, A., Foret, G., Meleux, F., Ung, A., Rouil, L., Ionescu, A., Candau, Y., and Beekmann, M.: Regional scale ozone data assimilation using an ensemble Kalman filter and the

- CHIMERE chemical transport model, *Geosci. Model Dev.*, 7, 283–302, doi:10.5194/gmd-7-283-2014, 2014.
- Grell, G. A., Peckham, S. E., Schmitz, R., McKeen, S. A., Frost, G., Skamrock, W. C., and Eder, B.: Fully coupled “online” chemistry in the WRF model, *Atmos. Environ.*, 39, 6957–6976, 2005.
- Guenther, A. B., Jiang, X., Heald, C. L., Sakulyanontvittaya, T., Duhl, T., Emmons, L. K., and Wang, X.: The Model of Emissions of Gases and Aerosols from Nature version 2.1 (MEGAN2.1): an extended and updated framework for modeling biogenic emissions, *Geosci. Model Dev.*, 5, 1471–1492, doi:10.5194/gmd-5-1471-2012, 2012.
- Herron-Thorpe, F. L., Mount, G. H., Emmons, L. K., Lamb, B. K., Chung, S. H., and Vaughan, J. K.: Regional air-quality forecasting for the Pacific Northwest using MOPITT/TERRA assimilated carbon monoxide MOZART-4 forecasts as a near real-time boundary condition, *Atmos. Chem. Phys.*, 12, 5603–5615, doi:10.5194/acp-12-5603-2012, 2012.
- Joiner J. and da Silva, A. B.: Efficient methods to assimilate remotely sensed data based on information content, *Q. J. Roy. Meteor. Soc.*, 124, 1669–1694, 1998.
- Klonecki, A., Pommier, M., Clerbaux, C., Ancellet, G., Cammas, J.-P., Coheur, P.-F., Cozic, A., Diskin, G. S., Hadji-Lazarou, J., Hauglustaine, D. A., Hurtmans, D., Khatatov, B., Lamarque, J.-F., Law, K. S., Nedelec, P., Paris, J.-D., Podolske, J. R., Prunet, P., Schlager, H., Szopa, S., and Turquety, S.: Assimilation of IASI satellite CO fields into a global chemistry transport model for validation against aircraft measurements, *Atmos. Chem. Phys.*, 12, 4493–4512, doi:10.5194/acp-12-4493-2012, 2012.
- Kukkonen, J., Olsson, T., Schultz, D. M., Baklanov, A., Klein, T., Miranda, A. I., Monteiro, A., Hirtl, M., Tarvainen, V., Boy, M., Peuch, V.-H., Poupkou, A., Kioutsioukis, I., Finardi, S., Sofiev, M., Sokhi, R., Lehtinen, K. E. J., Karatzas, K., San José, R., Astitha, M., Kallos, G., Schaap, M., Reimer, E., Jakobs, H., and Eben, K.: A review of operational, regional-scale, chemical weather forecasting models in Europe, *Atmos. Chem. Phys.*, 12, 1–87, doi:10.5194/acp-12-1-2012, 2012.
- MACC-II Final Report: Monitoring Atmospheric Composition and Climate – Interim Implementation, available at: <https://www.wmo.int/pages/prog/arep/gaw/documents/GAW-2013-Peuch-MACCII.pdf> (last access: 1 March 2016), 2014.
- Migliorini, S.: On the equivalence between radiance and retrieval assimilation, *Mon. Weather Rev.*, 140, 258–265, 2012.
- Migliorini, S., Piccolo, C., and Rodgers, C. D.: Use of the information content in satellite measurements for an efficient interface to data assimilation, *Mon. Weather Rev.*, 136, 2633–2650, 2008.
- Pagowski, M. and Grell, G. A.: Experiments with the assimilation of fine aerosols using an ensemble Kalman filter, *J. Geophys. Res.*, 117, D21302, doi:10.1029/2012JD018333, 2012.
- Pfister, G. G., Avise, J., Wiedinmyer, C., Edwards, D. P., Emmons, L. K., Diskin, G. D., Podolske, J., and Wisthaler, A.: CO source contribution analysis for California during ARCTAS-CARB, *Atmos. Chem. Phys.*, 11, 7515–7532, doi:10.5194/acp-11-7515-2011, 2011.
- Pfister, G. G., Walters, S., Emmons, L. K., Edwards, D. P., and Avise, J.: Quantifying the contribution of inflow on surface ozone over California during summer 2008, *J. Geophys. Res.-Atmos.*, 118, 12282–12299, doi:10.1002/2013JD020336, 2013.
- Rodgers, C. D.: *Inverse Methods for Atmospheric Sounding, Theory and Practice*, World Sci., Singapore, 2000.
- Smith, W. L. and Woolf, H.: The use of eigenvectors of statistical covariance matrices for interpreting satellite sounding radiometer observations, *J. Atmos. Sci.*, 33, 1–7, 1976.
- Thompson, O.: Regularizing the satellite temperature-retrieval problem through singular-value decomposition of the radiative transfer physics, *Mon. Weather Rev.*, 120, 2314–2328, 1992.
- Twomey, S.: Information content in remote sensing, *Appl. Optics*, 13, 942–945, 1974.
- Wiedinmyer, C., Akagi, S. K., Yokelson, R. J., Emmons, L. K., Al-Saadi, J. A., Orlando, J. J., and Soja, A. J.: The Fire INventory from NCAR (FINN): a high resolution global model to estimate the emissions from open burning, *Geosci. Model Dev.*, 4, 625–641, doi:10.5194/gmd-4-625-2011, 2011.
- Worden, H. M., Deeter, M. N., Edwards, D. P., Gille, J. C., Drummond, J. R., and Nedelec, P.: Observations of near-surface carbon monoxide from space using MOPITT multispectral retrievals, *J. Geophys. Res.*, 115, D18314, doi:10.1029/2010JD014242, 2010.