

A Hybrid MPI–OpenMP Parallel Algorithm and Performance Analysis for an Ensemble Square Root Filter Designed for Multiscale Observations

YUNHENG WANG AND YOUNGSUN JUNG

Center for Analysis and Prediction of Storms, University of Oklahoma, Norman, Oklahoma

TIMOTHY A. SUPINIE AND MING XUE

Center for Analysis and Prediction of Storms, and School of Meteorology, University of Oklahoma, Norman, Oklahoma

(Manuscript received 13 August 2012, in final form 16 January 2013)

ABSTRACT

A hybrid parallel scheme for the ensemble square root filter (EnSRF) suitable for parallel assimilation of multiscale observations, including those from dense observational networks such as those of radar, is developed based on the domain decomposition strategy. The scheme handles internode communication through a message passing interface (MPI) and the communication within shared-memory nodes via Open Multi-processing (OpenMP) threads. It also supports pure MPI and pure OpenMP modes. The parallel framework can accommodate high-volume remote-sensed radar (or satellite) observations as well as conventional observations that usually have larger covariance localization radii.

The performance of the parallel algorithm has been tested with simulated and real radar data. The parallel program shows good scalability in pure MPI and hybrid MPI–OpenMP modes, while pure OpenMP runs exhibit limited scalability on a symmetric shared-memory system. It is found that in MPI mode, better parallel performance is achieved with domain decomposition configurations in which the leading dimension of the state variable arrays is larger, because this configuration allows for more efficient memory access. Given a fixed amount of computing resources, the hybrid parallel mode is preferred to pure MPI mode on supercomputers with nodes containing shared-memory cores. The overall performance is also affected by factors such as the cache size, memory bandwidth, and the networking topology. Tests with a real data case with a large number of radars confirm that the parallel data assimilation can be done on a multicore supercomputer with a significant speedup compared to the serial data assimilation algorithm.

1. Introduction

With significant advances in computing power in recent years, advanced data assimilation (DA) techniques, such as the ensemble Kalman filter (EnKF) (Evensen 1994; Evensen and van Leeuwen 1996; Burgers et al. 1998; Houtekamer and Mitchell 1998; Anderson 2001; Bishop et al. 2001; Whitaker and Hamill 2002; Evensen 2003; Tippett et al. 2003) and four-dimensional variational data assimilation (4DVAR) (e.g., Le Dimet and Talagrand 1986; Courtier and Talagrand 1987; Sun and Crook 1997; Gao et al. 1998; Wu et al. 2000; Caya et al.

2005), are becoming more popular in both operational and research communities. However, they both incur a high computational cost, one of the biggest constraints for their operational applications at very high resolutions. Between EnKF and 4DVAR, the EnKF method appears to be more attractive for convective-scale numerical weather prediction (NWP), where nonlinear physical processes have critical roles. EnKF can also provide a natural set of initial conditions for ensemble forecasting. EnKF has been applied at scales ranging from global to convective and has produced encouraging results (e.g., Snyder and Zhang 2003; Dowell et al. 2004; Tong and Xue 2005, hereafter TX05; Xue et al. 2006; Jung et al. 2008; Buehner et al. 2010; Dowell et al. 2011; Hamill et al. 2011; Snook et al. 2011; Jung et al. 2012).

Among variants of EnKF, the ensemble square root Kalman filter (EnSRF) of Whitaker and Hamill (2002) is

Corresponding author address: Ming Xue, Center for Analysis and Prediction of Storms, University of Oklahoma, 120 David L. Boren Blvd., Norman, OK 73072.
E-mail: mxue@ou.edu

widely used in convective-scale DA studies involving radar data. The EnSRF, as well as the similar ensemble adjustment Kalman filter (EAKF; Anderson 2003) and the classic perturbed-observation EnKF algorithm (Evensen 2003), is an observation-space-based algorithm in which observations are assimilated one after another. Because of the sequential nature of the EnSRF (and EAKF and classic EnKF), parallelization of the algorithm at the observation level is not straightforward. It is possible to parallelize at the state variable level, that is, to perform the updating of the state variables in parallel because each observation updates many state variables within the covariance localization radius of the EnSRF, and these operations are independent. Such parallelization can be easily achieved on shared-memory platforms via Open Multiprocessing (OpenMP) directives, and is done with the Advanced Regional Prediction System (ARPS; Xue et al. 2003) EnSRF system (e.g., Xue et al. 2006; Jung et al. 2008). A processing element (PE) on a shared-memory or distributed-memory platform is an individual processor with single-core processors or a processor core on multicore CPUs. Each PE generally supports only a single process or a single thread. The number of PEs available on shared-memory nodes [the term “processing unit” (PU), will be used to refer to a shared-memory node] usually limits the scale of shared-memory parallelization (SMP) and the number of state variables that can be updated simultaneously. Distributed-memory parallelization (DMP) via the message passing interface (MPI) library would allow the use of much larger computers, which are essential for very high-resolution DA and NWP over large domains (Xue et al. 2007).

Anderson and Collins (2007, hereafter AC07) proposed a modification to the standard EAKF algorithm that is also applicable to EnSRF. In their algorithm, multiple observation priors (background converted to observed quantities via observation operators) are first calculated in parallel, and the observation priors corresponding to as-yet unused observations are updated by the filter together with the state vector, allowing easier parallelization at the state vector level (for a given observation, multiple elements in the state vector are updated in parallel). However, its state update procedure requires broadcasting the observation priors from one PU to the rest—and more importantly, the processing of observations is still serial. Because of this, the algorithm does not scale well when the number of PUs increases to the point where the cost of communication starts to dominate or when the ratio of the number of observations to that of state variables is large. Other parallel approaches have also been proposed by Keppenne and Rienecker (2002) and Zhang et al. (2005). While both

methods utilize domain decomposition, they differ in whether communication among PUs is allowed. Because there is no cross-PU communication in the algorithm of Zhang et al. (2005), the analysis near the PU boundaries is not the same as that of scalar implementation, which is a potentially serious drawback of their algorithm. Keppenne and Rienecker (2002), on the other hand, allow observations in other PUs to update the states in the current PU, but their communication cost is potentially very high because message passing is executed many times to properly exchange information among PUs.

In this paper, we develop a new parallelization algorithm for EnSRF (also suitable for other similar serial ensemble filters) that is especially suitable for dense observations that typically use relatively small horizontal covariance localization radii. Most NWP models, including the ARPS and the Weather Research and Forecasting (WRF) model, use horizontal domain decomposition for effective parallelization (Sathye et al. 1997; Michalakes et al. 2004). A domain-decomposition-based parallel DA strategy is attractive because it can share much of the parallelization infrastructure with the prediction model. If the DA system and prediction model use the same number and configuration of subdomains, then the transfer of model grids between the two systems will be more straightforward either through disk or within computer memory. Furthermore, with typical ensemble DA systems, the state arrays are usually moved between the prediction model and DA system through disk input/output (I/O) within the DA cycles; such I/O can take more than half of the total wall-clock time within each cycle (Szunyogh et al. 2008), making high-frequency assimilation of observations on large, high-resolution grids prohibitively expensive. Our eventual goal is to achieve data exchange through message passing within computer memory, bypassing disk I/O altogether; adopting a domain decomposition parallelization strategy would simplify this process. Finally, the domain decomposition strategy makes grid-based calculations within the DA system, such as spatial interpolation, easier.

The domain-decomposition-based strategy we propose takes advantage of the relatively small localization radii typically used by very dense observations within ensemble algorithms, because observations that do not influence state variables at the same grid points can be processed in parallel. More sparse conventional observations tend to require larger localization radii (Dong et al. 2011) and are therefore more difficult to process in parallel. In this case, a strategy similar to that of AC07 is taken, in which observations are processed serially but still using the same decomposed domains. Parallelization can be achieved at the state variable level in the case; in other words, different parallelization strategies

can be used in combination, taking advantage of the serial nature of the ensemble algorithms. Note that this approach scales well only for observations whose localization radius is large enough to impact most of the grid points in the model domain, unless additional steps are taken to balance the load, as in AC07.

In addition to domain-decomposition-based parallelization, we also want to take advantage of SMP capabilities of multicore compute nodes that are available on essentially all large parallel systems of today. SMP among cores on the same node eliminates explicit data transport among the cores, thus reducing communication costs and contention for interconnect ports. By performing domain decomposition for the nodes while parallelizing across the PEs (e.g., cores) on the same PUs (e.g., nodes), the decomposed domains can be larger relative to the localization radii, increasing the chance that observations on different decomposed domains can be processed independently.

For the EnSRF algorithm, SMP is easily achieved at the state variable level, because each observation will need to update all state variables within its localization radius, and these update operations are independent. Thus, the state variable update can be parallelized using OpenMP directives applied to the loops over the state variables. The combination of MPI and OpenMP strategies gives hybrid parallelization. This paper describes a hybrid parallel scheme implemented for the ARPS EnSRF system. In addition, observation data are organized into batches to improve the load balance when assimilating data from a number of radars.

This paper is organized as follows. Section 2 reviews the EnSRF formulation and briefly describes the ARPS model used in timing experiments. Section 3 introduces the parallel algorithms for high-density radar data and conventional observations separately. It also describes the OpenMP–MPI hybrid strategy as well as the observation organization. Validation of the parallel implementation and its performance are examined in section 4. A summary and conclusions are presented in section 5.

2. The ARPS ensemble DA system

The ARPS (Xue et al. 2000, 2001, 2003) model is a general purpose multiscale prediction system in the public domain. It has a nonhydrostatic, fully compressible dynamic core formulated in generalized terrain-following coordinates. It employs the domain decomposition strategy in the horizontal for massively parallel computers (Sathye et al. 1997; Xue et al. 2007), and it has been tested through real-time forecasts at convection-permitting and convection-resolving resolutions for many years (e.g., Xue et al. 1996), including forecasts in continental United

States (CONUS scale) domains at 4- and 1-km grid spacing (e.g., Xue et al. 2011), assimilating data from all radars in the Weather Surveillance Radar-1988 Doppler (WSR-88D) radar network using a 3DVAR method.

As mentioned earlier, the current ARPS EnKF DA system (Xue et al. 2006) is primarily based on the EnSRF algorithm of Whitaker and Hamill (2002). In addition, an asynchronous (Sakov et al. 2010) four-dimensional EnSRF (Wang et al. 2013) has also been implemented. The system includes capabilities for parameter estimation (Tong and Xue 2008), dual-polarimetric radar data assimilation (Jung et al. 2008), simultaneous reflectivity attenuation correction (Xue et al. 2009), and the ability to handle a variety of data sources (Dong et al. 2011). Additionally, it has been coupled with a double-moment microphysics scheme (Xue et al. 2010; Jung et al. 2012). To be able to apply this system to large convection-resolving domains, such as those used by ARPS 3DVAR for continental-scale applications (e.g., Xue et al. 2011), and to be able to assimilate frequent, high-volume observations, efficient parallelization of the system is essential.

Briefly, in EnSRF, the ensemble mean and ensemble deviations are updated separately. The analysis equations for the ensemble mean state vector $\bar{\mathbf{x}}$ and the ensemble deviations \mathbf{x}'_i are, respectively,

$$\bar{\mathbf{x}}^a = \bar{\mathbf{x}}^b + \boldsymbol{\rho}^o \mathbf{K} [\mathbf{y}^o - H(\bar{\mathbf{x}}^b)], \quad (1)$$

$$\mathbf{x}'_i{}^a = \beta(\mathbf{I} - \alpha \boldsymbol{\rho}^o \mathbf{K} H) \mathbf{x}'_i{}^b, \quad (2)$$

where \mathbf{K} is the Kalman gain and \mathbf{y}^o is the observation vector. Subscript i denotes the ensemble member and ranges from 1 to N , with N being the ensemble size; H is the forward observation operator that projects state variables to observed quantities, which can be nonlinear. Symbol \circ in the equations represents the Schur (elementwise) product and $\boldsymbol{\rho}$ is the localization matrix, containing localization coefficients that are typically functions of the distance between the observation being processed and the state variable being updated. The analysis background $\bar{\mathbf{x}}^b$ projected into observation space, that is, $H(\bar{\mathbf{x}}^b)$, is called the observation prior. Superscripts a , b , and o denote analysis, background, and observation, respectively. State vector \mathbf{x} includes in our case the gridpoint values of the three wind components (u , v , w), potential temperature (θ), pressure (p), the mixing ratios of water vapor (q_v), cloud water (q_c), rainwater (q_r), cloud ice (q_i), snow (q_s), and hail (q_h). When a two-moment microphysics parameterization scheme is used, the total number concentrations for the five water and ice species are also part of the state vector (Xue et al. 2010). Background state vectors $\bar{\mathbf{x}}^b$ and $\mathbf{x}'_i{}^b$ are either forecasts

from the previous assimilation cycle or the states updated by observations processed prior to the current one. The parameter β is the covariance inflation factor. Variable α is a factor in the square root algorithm derived by Whitaker and Hamill (2002):

$$\alpha = \left[1 + \sqrt{\mathbf{R}(\mathbf{H}\mathbf{P}^b\mathbf{H}^T + \mathbf{R})^{-1}} \right]^{-1}. \quad (3)$$

Here, \mathbf{R} is the observation error covariance matrix, \mathbf{P}^b is the background error covariance matrix, and \mathbf{H} is the linearized observation operator. The Kalman gain matrix \mathbf{K} is given by

$$\mathbf{K} = \mathbf{P}^b\mathbf{H}^T(\mathbf{H}\mathbf{P}^b\mathbf{H}^T + \mathbf{R})^{-1}. \quad (4)$$

In the above equation, matrices $\mathbf{P}^b\mathbf{H}^T$ and $\mathbf{H}\mathbf{P}^b\mathbf{H}^T$, representing the background error covariance between the state variables and observation priors, and that between observation priors, respectively, are estimated from the background ensemble, according to

$$\mathbf{P}^b\mathbf{H}^T = \frac{1}{N-1} \sum_{i=1}^N (\mathbf{x}_i^b - \bar{\mathbf{x}}^b) [H(\mathbf{x}_i^b) - \overline{H(\mathbf{x}^b)}]^T, \quad (5)$$

$$\mathbf{H}\mathbf{P}^b\mathbf{H}^T = \frac{1}{N-1} \sum_{i=1}^N [H(\mathbf{x}_i^b) - \overline{H(\mathbf{x}^b)}][H(\mathbf{x}_i^b) - \overline{H(\mathbf{x}^b)}]^T. \quad (6)$$

The overbars in Eqs. (5) and (6) denote the ensemble mean. When a single observation is analyzed, $\mathbf{P}^b\mathbf{H}^T$ becomes a vector having the length of the state vector \mathbf{x} . In practice, because of covariance localization, all elements in $\mathbf{P}^b\mathbf{H}^T$ are not calculated; those for grid points outside the localization radius of a given observation are assumed to be zero. In fact, it is this assumption that makes the design of our parallel algorithm practical; *observations whose domains of influence (as constrained by the covariance localization radii) do not overlap can be analyzed simultaneously*. Another basic assumption with this algorithm (and most atmospheric DA algorithms) is that observation errors are uncorrelated, so that observations can be analyzed sequentially in any order. When the observations are processed serially, one at a time, the observation error covariance matrix \mathbf{R} reduces to a scalar, as does matrix $\mathbf{H}\mathbf{P}^b\mathbf{H}^T$. In this case, $\mathbf{H}\mathbf{P}^b\mathbf{H}^T$ is the background error variance at the observation point.

After an observation is analyzed based on Eqs. (1)–(6), the analyzed ensemble states \mathbf{x}_i^a ($i = 1, \dots, N$), the sum of the ensemble mean and deviations, become the new background states \mathbf{x}_i^b for the next observation, and the analysis is repeated until all observations at a given

time are analyzed. An ensemble of forecasts then proceeds from the analysis ensemble until the time of new observation(s); at that time the analysis cycle is repeated.

3. The parallel algorithm for EnSRF

For convective-scale weather, Doppler weather radar is one of the most important observing platforms. The U.S. National Weather Service (NWS) operates a network of over 150 WSR-88D radars that continuously scan the atmosphere, at a rate of one full volume scan every 5–10 min, producing radial velocity and reflectivity data. One volume scan in precipitation mode typically contains 14 elevations with approximately several million observations every 5 min.

The number of conventional observations, such as surface station measurements, upper-air soundings, and wind profiler winds, is small compared to radar observations; because the observations typically represent weather phenomena of larger scales, their assimilation in EnKF typically uses larger covariance localization radii, and therefore their influence reaches larger distances (Dong et al. 2011). Because of the different characteristics of each data type, different parallel strategies are employed for conventional and radar data.

a. The parallel algorithm for high-density observations with small covariance localization radii

The algorithm partitions the entire analysis domain into subdomains defined by the number of participating MPI processes in the horizontal x and y directions. No decomposition is performed in the vertical direction; therefore, state variables are always complete in the vertical columns. High-density radar observations (and other high-resolution observations including those of satellite) are distributed to each subdomain according to their physical locations. Figure 1 illustrates an analysis domain that is partitioned into four physical subdomains horizontally, to be handled by four PUs in the computing system. Each computational domain is composed of the physical subdomain (in darker gray for P1, separated with thick solid lines) and extended boundary “halo” zones surrounding the physical subdomain (in light gray for P1, bounded by thin lines); the physical domain and the boundary halo zones combined together are called computational subdomains. The width of the extended boundary halo zone for the DA system is typically larger than the halo zone or “ghost cells” needed for boundary condition exchanges in parallel NWP models based on domain decomposition (e.g., Sathye et al. 1997). The width of the halo zone in the ARPS model, for example, is only one grid interval on each boundary.




FIG. 1. A schematic of the domain decomposition strategy for the analysis of high-density observations, illustrated with four PUs (denoted by P1–P4). Letters i – l denote observations that are assumed to be equally spaced, and letters a – h indicate the influence limits (as determined by the covariance localization radii of EnKF) of those observations. In this example, observations i and l are far enough apart that they will not influence any of the same state variables; they are among the observations that are analyzed simultaneously in the first step of the procedure. Observations j and k are analyzed in the second step, but they must be analyzed sequentially. Note that in practice, there will be many more observations within patches S1 and S2 of subdomains P1–P4 than shown in the figure.

The extended boundary zone on each side must be at least as wide as the maximum localization radius (R) of observations handled by the algorithm in the subdomain. For radar observations, R is usually equal to a few grid intervals. Each physical subdomain is further divided into four patches that are separated by bold dashed lines in Fig. 1, and these patches are labeled S1, S2, S3 and S4, respectively. The horizontal width of patch S2 and the vertical height of patch S3 must be at least $2R$. The rest of the physical domain is assigned to patches S1 and S4 as in Fig. 1, and their horizontal width and height also must be at least $2R$. Thus, the width of the physical subdomain must be larger than $4R$ for the algorithm to work. All other subdomains in Fig. 1 are divided following the same patch pattern. Such a patch division assures that patches with the same label in adjacent subdomains are at least $2R$ apart, so observations in any one patch do not affect grid points in the same patch on other PUs; thus, they can be analyzed in parallel. In other words, no two observations that are being analyzed in parallel will influence the same grid point. In

practice, we want to make patch S1 as large as possible, increasing the chance that any two observations can be processed independently (see below). Thus, the width of S2 and the height of S3 are assigned the minimum possible size of $2R$ (see Fig. 1), which leaves the majority of the subdomain to patch S1.

The EnKF DA over the analysis domain is performed in four sequential steps for observations within S1–S4. In the first step, only observations within S1 on all PUs are assimilated in parallel, while observations on each S1 patch are assimilated sequentially. Let P be the number of PUs. Then, there can be at most P observations being assimilated in parallel at any time. After all observations located within S1 are assimilated, MPI communications are required to properly update state variables at grid points within the extended boundary zones that are shared with neighboring PUs. The same procedure is then repeated for observations within S2–S4 in steps 2–4.

The assimilation of observations within the same-labeled patches from all PUs can be done in parallel because 1) the grid points influenced by the observations analyzed in parallel are separated far enough without overlap; and 2) the ensemble state arrays are extended beyond the physical subdomain, so that the influence on state grids by observations within each subdomain can be passed to its neighbor PUs with MPI communications. Best load balancing is realized if the same-labeled patches contain the same number of observations, so that all PUs can complete each analysis step in approximately the same time. In practice, however, the number of observations on each subdomain is usually different because of uneven spatial distribution of observations (and of observation types). One way to improve parallelism is to make one patch (S1 in our system) as large as possible, which increases the number of observations that can be processed independently and improves the load balance. Assimilation of observations on S2–S4 may not be well balanced. However, because they tend to be smaller and contain fewer observations, their effect on the assimilation time tends to be small.

Since high-density observations, such as radar data, usually assume relatively small localization radii, the constraint that the width of the physical subdomain should be at least $4R$ in each direction usually does not become a major problem, especially when the DA domain is large. When a hybrid MPI–OpenMP parallelization strategy is used, this problem can be further alleviated (see later). While the proposed algorithm is valid for most meteorological observations that can assume a small localization radius, certain “integral observations,” such as radar reflectivity with path-integrated attenuation effect (e.g., Xue et al. 2009) and GPS slant-path water vapor (e.g., Liu and Xue 2006),

pose special challenges for the serial EnSRF algorithm in general, since their observation operators are non-local (Campbell et al. 2010).

b. The parallel algorithm for conventional observations with large covariance localization radii

Currently supported conventional observations in the ARPS EnKF system include surface station, upper-air sounding, wind profiler, and aircraft observations. Since the covariance localization radii applied to these observations are usually large, the width of the extended boundary zones described in section 3a would be impractical for these data, unless the decomposed subdomains are much larger than the localization radii. This is usually only true when a small number of subdomains is used. Therefore, we design and implement an alternative algorithm for this type of observations. Because the number of conventional (or any other coarse resolution) observations is typically much smaller than the number of (dense) radar observations, we can afford to process the observations serially while trying to achieve parallelism at the state variable level, similar to the strategy taken by AC07.

In our current implementation, conventional observations within the entire analysis domain are broadcast to all PUs and assimilated one by one. Only the PU containing the observation to be analyzed computes the observation prior; it then broadcasts the observation prior ensemble, $H(\mathbf{x}_i)$, to all other PUs. The state variables within the covariance localization radius of this observation are updated simultaneously on each PU that carries the state variables (Fig. 2). Since we do not need extra boundary zones, state variable updating occurs within the computational subdomains of the original NWP model. However, a set of MPI communications between PUs is still needed right after the analysis of each observation to update the state variables within the halo zone to facilitate the spatial interpolation involved in observation operators. These steps are repeated until all observations are assimilated.

Our current implementation does not precalculate or update $H(\mathbf{x})$ as part of the extended state vector as AC07 does, and we use a regular domain decomposition strategy to distribute the state variables across the PUs. This implementation will have load balance issues for conventional observations, especially when the covariance localization radii of these observations are small relative to the size of the entire model domain. AC07 mitigates this problem by distributing the state variables across PUs as heterogeneously as possible, that is, by distributing neighboring grid points across as many PUs as possible. Such an irregular distribution of state variables makes it




FIG. 2. A schematic for analyzing conventional data. Three steps are involved when analyzing one observation whose location is denoted by a black dot in the figure: 1) PU14 computes $H(\mathbf{x}_i)$ (where i is the ensemble index), 2) $H(\mathbf{x}_i)$ are broadcasted to all PUs, and 3) state variables \mathbf{x}_i within the influence range of this observation (within the large circle) are updated in parallel by the PUs that carry the state variables.

difficult to implement gridpoint-based treatments within the EnKF algorithms. The $H(\mathbf{x})$ precalculation and update strategy employed by AC07 allows simultaneous calculation of observation priors. This can be an option in a future implementation; in fact, the 4D EnSRF algorithm implemented by Wang et al. (2013) employs this strategy.

c. Hybrid MPI–OpenMP parallelization

All current supercomputers use compute nodes with multiple shared-memory cores. The original ARPS EnSRF code supports OpenMP parallelization via explicit loop-level directives at the state variable update level (Xue et al. 2006). Thus, it is straightforward to employ a hybrid technique, using SMP among cores on the same node and DMP via MPI across nodes. Doing so can reduce explicit data communication within nodes and allow for larger S1 patches within the decomposed domains on each PU (see Fig. 1). Our hybrid implementation is designed such that each MPI process spawns multiple threads. Since message passing calls are outside of the OpenMP parallel sections, they are parallel thread safe, that is, only the master thread in a process makes calls to MPI routines. The final program is flexible enough to run in MPI only, OpenMP only, or in MPI–OpenMP hybrid modes, on a single-node workstation or on supercomputers made up of multiple nodes.

d. Parallel strategy for assimilating data from multiple radars

In the ARPS EnKF system, full-resolution radar observations in the radar coordinates are usually mapped horizontally to the model grid columns during

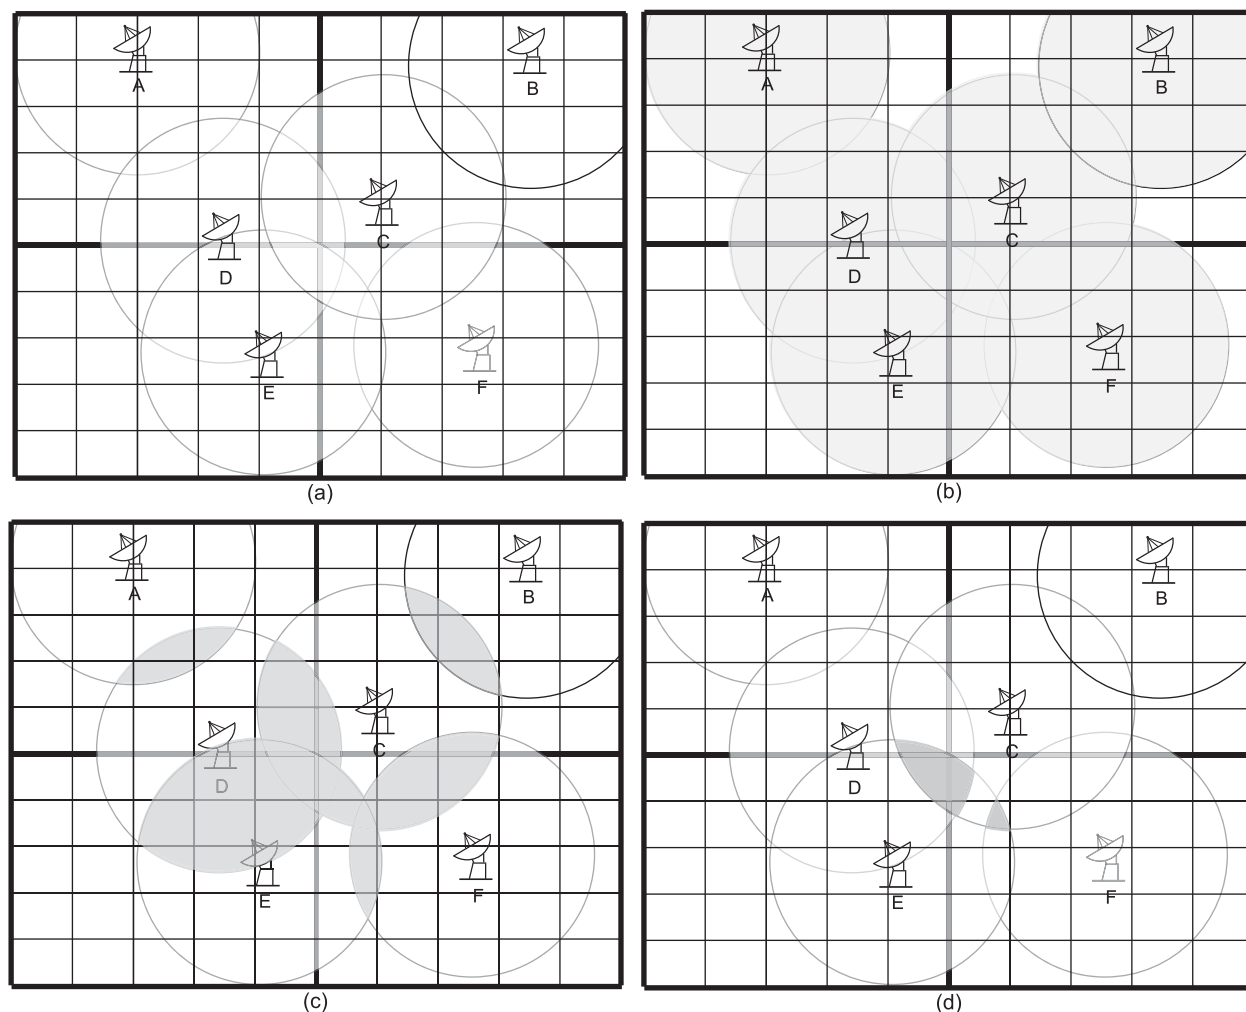


FIG. 3. Composite radar data batches organized such that within each batch, no more than one column of data exists for each grid column. (a) Observations from six radars (A–F) with their coverage indicated by the maximum range circles are remapped onto the model grid. (b) Observations of the first batch. (c) Observations of the second batch. (d) Observations of the third batch. If there are more observations unaccounted for, then additional data batch(es) will be formed.

preprocessing (Brewster et al. 2005). The original ARPS EnSRF implementation processes data from one radar at a time, sequentially. This is convenient because the data are stored in arrays for individual radars on elevation levels (Xue et al. 2006). For data from the same radar, only a few parameters are needed to describe the radar characteristics. However, because each radar typically covers only a portion of the model domain, this procedure severely limits the scalability of the analysis system because of load imbalances (see Fig. 3). Figure 3a illustrates a domain that contains six radars labeled A–F. If this domain is decomposed into four subdomains, then all PUs, except P1, will be idle when data from radar A are assimilated. The same is true for radars B–F. To mitigate this problem, we develop a procedure that merges radar data into composite sets or batches so that

data from multiple radars can be processed at the same time.

In the analysis program, all vertical levels of radar observations at each horizontal grid location are stored continuously as a vector column. The most general approach is to store all columns of radar data in a single dynamically allocated storage array or data structure while keeping track of the radar characteristics associated with each column. Each column may contain different numbers of available radar elevations. When overlapping coverage exists, the grid columns covered by multiple radars will have multiple columns of data (see Fig. 3a). To keep track of data in reference to the analysis grid, it is convenient to define arrays that have the same dimensions as the model grid in the horizontal directions, but such arrays will only be able to store no

more than one column of data at each grid location unless the last dimension is defined dynamically or predefined to be large enough. While for optimally tuned EnKF the order in which observations are assimilated should not matter, in practice, because the ensemble spread can be reduced too much by observations processed earlier before covariance inflation is applied, the order of observation processing sometimes does matter somewhat. For this reason, we group the radar data into several batches, the number of which is no bigger than the maximum number of radars covering the same spot anywhere in the analysis domain. For a radar network that is designed to maximize spatial coverage, such as the WSR-88D radar network, this maximum is usually a single digit number; that is, anywhere in the network, fewer than 10 radars observe the same column.

Figure 3 shows the spatial coverage of three batches of data that add up to all the columns of data available; those three batches of observations will be processed in sequence. Within regions having multiple radar coverage, the radar from which data will be first picked can be chosen randomly or based on the order the data were input into the program. Alternatively, the data columns from the closest radar can be picked first. The last option is more desirable, as it removes the randomness of the algorithm. Finally, because the radar data are no longer organized according to radar, additional two-dimensional arrays are needed to store parameters for each data column. When only a few elevations within a radar volume scan are analyzed using short (e.g., 1–2 min) assimilation cycles, the vertical dimension of the arrays storing the composite datasets need only to be a few.

With the above-mentioned implementation, the load balance is significantly improved for the first composite dataset. It should be noted that we usually assimilate reflectivity data even in precipitation-free regions, which has the benefit of suppressing spurious storms (TX05). We note that load imbalance does still exist with radial velocity data in the first group, since they are usually only available in precipitation regions; however, their numbers are usually much smaller than the total number of reflectivity data. In addition, load imbalances usually exist with the second group of data and above, but again the volume of data in these groups is small since they only exist in overlapping regions, and these regions are usually spread over the assimilation domain.

4. Algorithm verification and performance analysis

a. Verification of the parallelized code

The domain partition and batch processing inevitably change the sequence of observations being assimilated

in the EnKF system. Theoretically, the order in which the observations are processed does not matter for observations with uncorrelated errors, to the extent that sampling error does not impact the results. In practice, the analysis results may differ significantly if the filter is not properly tuned, where the tuning typically includes covariance inflation and localization.

A set of experiments has been performed to investigate the effect of domain decomposition on the analysis of simulated radar observations in an observing system simulation experiment (OSSE) framework. Convective storms are triggered by five 4-K ellipsoidal thermal bubbles with a 60-km horizontal radius and a 4-km vertical radius in an environment defined by the 20 May 1977 Del City, Oklahoma, supercell sounding (Ray et al. 1981). The model domain is $300 \times 200 \times 16 \text{ km}^3$ with horizontal and vertical grid spacings of 1 km and 500 m, respectively. Forty ensemble members are initiated at 3000 s of model time. The full state vector has 1.4×10^9 elements. Simulated radar observations from three radars are produced, using the standard WSR-88D volume coverage pattern (VCP) 11, which contains 14 elevation levels. The total number of observations is approximately 6.7×10^5 from three volume scans spanning 5 min each. Radar DA is first performed at 5-min intervals from 3300 to 5700 s, using the original serial ARPS EnSRF code to provide an ensemble for subsequent parallel assimilation tests. The Milbrandt and Yau (2005) double-moment microphysics scheme is used in both truth simulation and DA. The environment and model configurations that are not described here can be found in Xue et al. (2010).

Three parallel DA experiments are then performed at 6000 s, one running in pure MPI mode, one in pure OpenMP mode, and one in pure OpenMP mode but processing observations serially in a reversed order. These experiments are referred to as MPI, OMP_F, and OMP_B (F for forward and B for backward), respectively. For each experiment, average RMS errors for the state variables are computed against the truth simulation at the grid points where truth reflectivity is greater than 10 dBZ. The RMS errors of MPI and OMP_B are normalized by the RMS errors of OMP_F and shown in Fig. 4 for individual state variables. Most of the normalized errors are very close to 1, and all of them are between 0.95 and 1.05 for MPI. Among the variables, the total number concentration for rainwater shows the largest variability, probably because of the high sensitivity of reflectivity to the raindrop size distribution. In fact, the normalized error for rainwater number concentration is an outlier for OMP_B, reaching close to 1.25, much larger than the normalized error of about 1.05 for MPI. These results suggest that the effect of the domain partition on the analysis is small, and the differences are




FIG. 4. RMS errors averaged over the grid points where truth reflectivity is >10 dBZ and normalized by the errors of experiment OMP_F. The state variables are the 16 ARPS prognostic variables (refer to text) and their respective number concentrations (N_{tc} , N_{tr} , N_{ti} , N_{ts} , and N_{th} , associated with a two-moment microphysics scheme used).

within the range of sampling uncertainties of the ensemble system.

With respect to the parallel code implementation for conventional data analysis, domain decomposition does not change the sequence of the observation processing (see section 3b). Therefore, identical results from experiments OMP_F and MPI are guaranteed. The results from the experiments when simulated surface observations are also included are not shown here.

b. Performance evaluation with OSSE experiments

The performance of our parallel EnKF system is evaluated with radar DA benchmark experiments on a Cray XT5 system (called Kraken) at the National Institute for Computational Sciences (NICS) at the University of Tennessee, which has 9408 total compute nodes with 12 cores each (6 cores per processor, 2 processors per node), giving a peak performance of 1.17 petaflops. With Kraken, users can set the number of MPI processes per node (1–12), the number of MPI processes per processor (1–6), and the number of cores (OpenMP threads) per MPI process (1–12). A number of experiments with combinations of different numbers of MPI processes, OpenMP threads, cores per node, and cores per processor have been performed to examine the timing performance of various configurations. The same case described in section 4a is used for benchmarking.

First, the scalability of the OpenMP implementation is investigated as a reference. Since each Kraken node

contains only 12 cores, the maximum number of threads that can be used for an OpenMP job is 12. The OpenMP implementation shows scalability up to 8 cores (see Table 1), beyond which the reduction in wall-clock time becomes minimal. One very likely reason is the contention accessing shared memory and cache by different cores of the Opteron processors used.

To evaluate the performance of our MPI implementation, we ran several OpenMP and MPI experiments on

TABLE 1. Timing comparisons of OpenMP experiments with MPI experiments on one compute node. Speedup for OpenMP and MPI experiments are computed relative to o1 and m01 \times 01, respectively.

Experiment	Total number of cores used	Wall-clock time (s)	Speedup
o1	1	6310	1.00
o2	2	3617	1.75
o4	4	2597	2.43
o6	6	1919	3.29
o8	8	1597	3.95
o12	12	1607	3.93
m01 \times 01	1	6815	1.00
m01 \times 02	2	3994	1.71
m02 \times 01	2	5698	1.20
m01 \times 04	4	2660	2.56
m02 \times 02	4	3690	1.85
m04 \times 01	4	2896	2.35
m03 \times 02	6	4177	1.63
m02 \times 04	8	2100	3.25
m04 \times 02	8	2413	2.82

TABLE 2. Timing comparisons of pure MPI experiments with hybrid MPI–OpenMP experiments on four compute nodes. Speedup is computed relative to experiment m01 \times 01 (6815 s in Table 1).

Experiment	Total number of cores used	Wall-clock time (s)	Speedup
m01 \times 04_01	4	2343	2.91
m02 \times 02_01		3577	1.91
m04 \times 01_01		2750	2.48
m02 \times 04_02	8	2169	3.14
m04 \times 02_02		2330	2.92
m03 \times 04_03	12	1575	4.33
m06 \times 02_03		1699	4.01
m04 \times 04_04	16	1327	5.14
m02 \times 10_05	20	915	7.45
m10 \times 02_05		1357	5.02
m04 \times 05_05		1082	6.30
m05 \times 04_05		1085	6.28
m03 \times 08_06	24	880	7.74
m06 \times 04_06		1049	6.50
m04 \times 10_10	40	637	10.70
m10 \times 04_10		720	9.47
m06 \times 08_12	48	606	11.25
h01 \times 04_01o2	8	1471	4.63
h01 \times 04_01o4	16	1129	6.04
h01 \times 04_01o6	24	831	8.20
h01 \times 04_01o8	32	772	8.83
h01 \times 04_01o12	48	733	9.30
h02 \times 04_02o2	16	1200	5.68
h02 \times 04_02o4	32	908	7.51
h02 \times 04_02o6	48	709	9.61

a single compute node. Table 1 lists the wall-clock times and relative speedups for these experiments. The experiment names follow the convention o(total cores used) for OpenMP and m(nproc_x) \times (nproc_y) for MPI experiments, where nproc_x and nproc_y denote the number of PUs corresponding to the decomposed domains in the x and y directions, respectively. Generally, the OpenMP jobs perform better than their MPI counterparts using the same number of cores when running on a single node because of the communication overhead of MPI processes and possibly better load balance with OpenMP. It is also noticed that the wall-clock time is heavily influenced by the domain partitioning configuration in the x and y directions. For example, m02 \times 01 takes almost 1.4 times longer than m01 \times 02, although both use the same number of cores. Since FORTRAN arrays are stored contiguously in the column-major order in the computer memory, a run that has a smaller partition number in the x direction than the y direction (e.g., m01 \times 02) is better at taking advantage of the spatial

locality of the data in memory. This can accelerate data loading from main memory into cache and improve cache reuse. Conversely, an inefficient partition can degrade the performance even when more system resources are used. For example, m03 \times 02 using six cores has a much smaller speed improvement over m01 \times 01 than experiments using four cores or even some experiments using two cores. These results suggest that finding the optimal domain decomposition is important in achieving the best performance with the given system resources.

Table 2 shows performance data collected from pure MPI runs, and from hybrid MPI–OpenMP experiments that run on four Kraken nodes. All experiments are named as follows: m(h)(nproc_x) \times (nproc_y)_(number of processes per node) o(number of threads per process), where m denotes MPI runs and h denotes hybrid runs. For MPI runs, the number of threads per process is always 1. Thus, o(number of threads per process) is omitted from the notations for all MPI runs in Table 2. Since each Kraken node contains two processors, the processes on each node are distributed to those processors as evenly as possible in order to obtain the best possible performance.

It is found that the domain partitioning again plays an important role in the DA system performance. For example, experiments that use 20 cores in total on four compute nodes show large variability in the execution time. Among these experiments, m02 \times 10_05 has the best performance, suggesting that m02 \times 10_05 utilizes the system cache most efficiently and/or has the least message-passing overhead given 20 cores. Generally, the MPI experiments using more nodes perform better than those experiments with the same domain partitioning but using fewer nodes. For example, m01 \times 04 in Table 1 running on one compute node takes 2660 s to finish, while m01 \times 04_01 in Table 2 running on four compute nodes takes only 2343 s. This is consistent with the observation that performance is improved as available cache size increases. Adding more processes improves the performance on four compute nodes. As an example, m06 \times 08_12 takes less time than those experiments using 40 cores or less. This is because more observations can be processed in parallel in the m06 \times 08_12 experiment than the others, even though MPI communication costs are higher than in the other experiments. However, as observed before with OpenMP experiments, access contention for the memory bandwidth and the cache sharing as more cores are used may impede the performance at some point. It suggests that there is a trade-off between the number of processes and available computing resources and, therefore, finding optimal configurations for MPI runs may not be straightforward because it depends on a number of hardware factors.

TABLE 3. Comparison of the minimum time taken in hybrid mode with that in MPI mode using the same number of cores on four compute nodes.

No. of cores	Hybrid case	Min time (s)	MPI case	Min time (s)	Difference (s)
8	$h01 \times 04_01o2$	1471	$m02 \times 04_02$	2169	698
16	$h01 \times 04_01o4$	1129	$m04 \times 04_04$	1327	198
24	$h01 \times 04_01o6$	831	$m03 \times 08_06$	880	49
40	$h02 \times 10_05o2$	635	$m04 \times 10_10$	637	2
48	$h03 \times 08_06o2$	604	$m06 \times 08_12$	606	2

For the hybrid runs, the wall-clock times of $m01 \times 04_01$ (i.e., $h01 \times 04_01o1$), $h01 \times 04_01o2$, $h01 \times 04_01o4$, $h01 \times 04_01o6$, $h01 \times 04_01o8$, and $h01 \times 04_01o12$ decrease monotonically, in that order. The decreasing trend of wall-clock time with an increasing number of threads is consistently found in other similar sets of experiments. It is also found that the hybrid runs are as sensitive as the MPI runs to the domain partitioning, available cache, and other hardware configuration factors. A hybrid experiment can outperform or underperform the corresponding MPI experiments using the same resources (number of cores and number of nodes) depending on the configuration (Tables 2 and 3). For example, the minimum wall-clock time with eight cores from four nodes in hybrid mode is 1471 s, which is smaller than the minimum time required by an MPI run with eight processes on four nodes (2169 s) in Table 3. On the other hand, $h01 \times 04_01o12$ takes 733 s, more than the 606 s of $m06 \times 08_12$, which uses the same resources. It is also observed that a larger improvement is achieved by the hybrid jobs with a fewer number of threads. This is because observations are processed one by one with OpenMP processes. By using more MPI processes rather than more OpenMP threads, we can assimilate more observations simultaneously and, hence, improve the

parallel efficiency (see section 4c for more details). In addition, cache availability and memory access contention with a large number of threads in the hybrid experiments also affect program performance.

c. Performance evaluation with a real data application

The parallel ARPS EnKF system is applied to the 10 May 2010 Oklahoma–Kansas tornado outbreak case. Over 60 tornadoes, with up to EF4 intensity, affected large parts of Oklahoma and adjacent parts of southern Kansas, southwestern Missouri, and western Arkansas on that day. This real data case is run on an SGI UV 1000 cache-coherent (cc) nonuniform memory access (NUMA) shared-memory system at the Pittsburgh Supercomputing Center (PSC). The system, called Blacklight, is composed of 256 nodes containing 2 eight-core Intel Xeon processors each; its theoretical peak performance is 37 teraflops. The cc-NUMA architecture allows for SMP across nodes. Up to 16 terabytes (TB) of memory can be requested for a single shared-memory job, while hybrid jobs can access the full 32 TB of system memory.

The EnSRF analyses are performed on a grid with 4-km horizontal grid spacing, using 40 ensemble members. The domain consists of $443 \times 483 \times 53$ grid points,

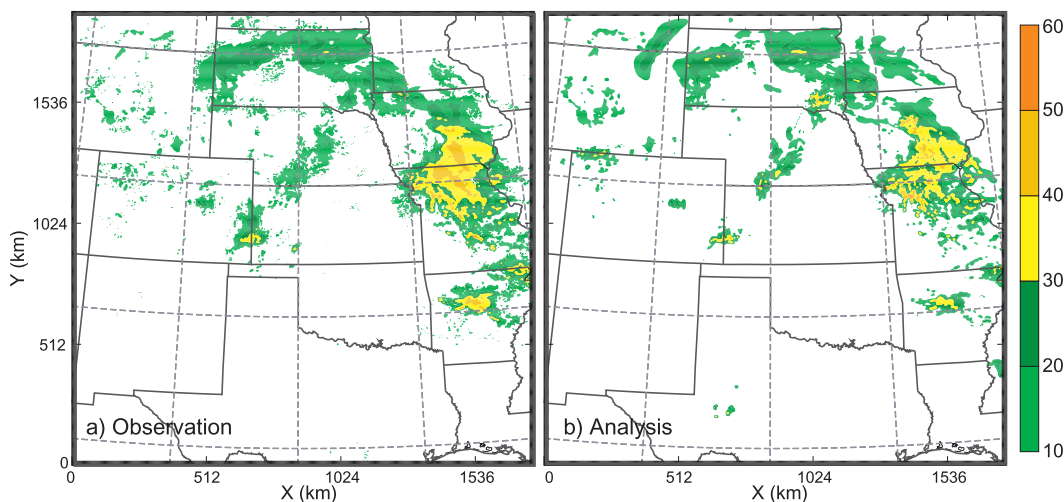


FIG. 5. (a) The observed radar reflectivity mosaic and (b) reflectivity field analyzed by the parallel EnKF algorithm, at model grid level 20 at 1800 UTC 10 May 2010.

TABLE 4. Names and configurations of real data experiments.

Experiment		Number of PUs in x direction	Number of PUs in y direction	Number of threads per PU	Total number of cores used
OpenMP	o16			16	16
	o32			32	32
	o64			64	64
	o80			80	80
	o160			160	160
MPI	m16	1	16		16
	m32	2	16		32
	m64	3	16		64
	m80	5	16		80
	m160	10	16		160
Hybrid group 1	h4o4	1	4	4	16
	h8o4	1	8		32
	h16o4	2	8		64
	h20o4	2	10		80
	h40o4	4	10		160
Hybrid group 2	h2o8	1	2	8	16
	h4o8	1	4		32
	h8o8	1	8		64
	h10o8	2	5		80
	h20o8	4	5		160
Hybrid group 3	h2o16	1	2	16	32
	h4o16	1	4		64
	h5o16	1	5		80
	h10o16	2	5		160

node. Tests with split files on this system, corresponding to $h04 \times 08_01o8$ (see above-mentioned naming conventions), reveal that the times spent on I/O and message passing are reduced (the latter because of the reduced exchanges of gridded information across MPI processes); the total wall-clock time for I/O and message passing for one experiment was reduced from 1231 to 188 s using split files.

5. Summary and conclusions

A parallel algorithm based on the domain decomposition strategy has been developed and implemented within the ARPS EnKF framework. The algorithm takes advantage of the relatively small spatial covariance localization radii typically used by high-resolution observations such as those of radar. Assuming that the maximum horizontal covariance localization radius of the observations to be analyzed in parallel is R , the horizontal area of a decomposed physical subdomain should be at least $4R \times 4R$. An additional boundary zone of width R is added to each side of the physical subdomains to create enlarged computational subdomains, which facilitate information exchanges between neighboring subdomains. Each subdomain is assigned to one processing unit (PU), within which no MPI message passing is required. The

subdomains are then further divided up into four subpatches, denoted S1–S4. The width and height of each patch are required to be at least $2R$ to ensure any two observations that may be processed in parallel are well separated. In practice, the size of S1 is made as large as possible within its subdomain to increase the probability that observations from different subdomains can be processed in parallel.

Observations within the four patches are processed sequentially, but data in the patches with the same label in different subdomains are processed simultaneously. Distributed-memory parallelization is therefore achieved at the observation level. The patch division ensures that most of the analysis work is done in parallel when processing observations within patches S1 of all PUs. To handle the load imbalance issue when assimilating observations from many radars, the observation arrays are organized into batches. The maximum number of batches is limited by the maximum number of radars covering the same location anywhere in the analysis domain. Such an observation organization improves the uniformity of observation distribution within the first observation batch and thereby improves load balance.

Conventional data that use larger covariance localization radii are still processed serially. State variables

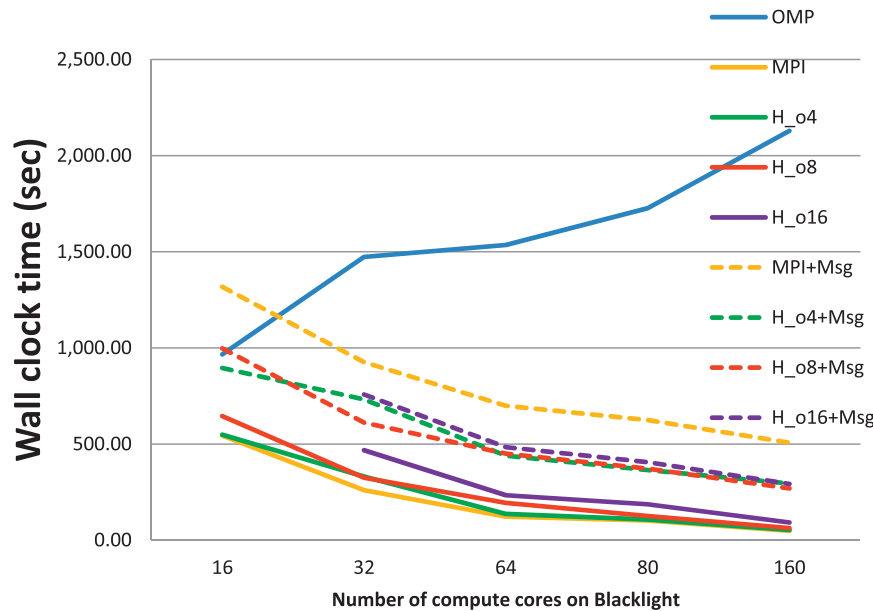


FIG. 7. Wall-clock times of the EnKF analyses as a function of the total number of compute cores used, for the 10 May 2010 real data case in the analysis domain shown in Fig. 6, obtained on the PSC Blacklight (an SGI UV 1000). Hybrid runs with 4, 8, and 16 OpenMP threads within each MPI process are denoted as H_o4, H_o8, and H_o16, respectively. In all cases, all cores on the compute nodes were fully utilized, either by individual MPI processes or by OMP threads. Solid lines denote the total time excluding message passing, and dashed lines show the total times including message passing. Data I/O times are excluded from all statistics.

influenced by a particular observation are updated synchronously on the PUs carrying those state variables.

The algorithm supports three parallel modes: pure OpenMP, pure MPI, and MPI–OpenMP hybrid. Within the PUs with multiple cores, shared-memory parallelization can be achieved via OpenMP at the state variable update level. OpenMP parallelization reduces message passing overhead and allows for larger decomposed domains, making the 4R requirement easier to satisfy.

It was first confirmed via OSSEs that changing the sequence of observation processing because of domain decomposition has little impact on the analysis. Parallel DA benchmark experiments are performed on a Cray XT5 machine. The OpenMP implementation shows scalability up to eight threads, beyond which memory and cache access contention limit further improvement. MPI and OpenMP runs on a single compute node show that OpenMP parallelization runs faster because of the lower communication overhead. MPI jobs with a smaller number of partitions in the x direction than in the y direction exhibit better performance. The same also applies to most of the hybrid jobs, although all hybrid jobs do not outperform the corresponding MPI jobs.

A real data case involving 35 radars is tested on an SGI UV 1000 cc-NUMA system capable of shared-memory programming across physical nodes. Poor scalability with

pure OpenMP is observed when more than one node is used, but both MPI and hybrid runs show good scalability on this system. Excluding message passing time, pure MPI runs exhibit the best performance. When message-passing time is included, the hybrid runs generally outperform pure MPI runs. For this real data case, the EnKF analysis excluding I/O can be completed within 4.5 min using 160 cores of the SGI UV 1000.

Given a fixed amount of resources, the hybrid jobs improve more over pure MPI jobs with fewer numbers of threads. Because MPI processes realize parallelization at the observation level, they are more efficient than OpenMP threads. However, there is a trade-off between a performance improvement because of the parallel processing of observations and degradation because of increased message passing overhead. On the other hand, a pure OpenMP strategy for EnKF shows good scalability on symmetric shared-memory systems but is limited by the number of cores available on the individual node and by the physical memory available on the node. With pure OpenMP, data I/O can only be handled by a single process, reducing the overall scalability.

The MPI–OpenMP hybrid strategy combines the strengths of both methods. However, care must be taken when partitioning the domain, because the configuration of MPI domain partitioning has a significant impact on

the system performance. Given the same resources, jobs with smaller numbers of partitions in the x direction tend to run faster because FORTRAN arrays are stored in the column-major order in memory. Timing experiments have also shown that determining the optimal decomposition configuration on a specific computing system is not straightforward because the performance depends on factors such as the subdomain size in the x and y directions, the number of cores on each node, the cache sizes and memory bandwidth available to each core, and the networking topology across the nodes.

In all configurations, data I/O constituted a large portion of the execution time. Experiments on a small dedicated Linux cluster show that the time spent on I/O and message passing are reduced significantly by distributing I/O loads among the MPI processes with MPI–OpenMP hybrid or pure MPI runs.

Although a data batching strategy is developed to reduce the load imbalance issue, further improvement could be obtained through dynamic load balancing. Another problem is the low resource utilization during internode communications because all threads are idle except one: the master thread. The development of runtime management algorithms, for example, the Scalable Adaptive Computational Toolkit (SACT) (Li and Parashar 2007; Parashar et al. 2010), are expected to decrease runtime of the application automatically with reduced efforts from developers. Finally, we point out that our parallel algorithm can be easily applied to other serial ensemble-based algorithms such as EAKF and the classic EnKF.

Acknowledgments. This work was primarily supported by NSF Grants OCI-0905040 and AGS-0802888, and NOAA Grant NA080AR4320904 as part of the Warn-on-Forecast Project. Partial support was also provided by NSF Grants AGS-0750790, AGS-0941491, AGS-1046171, and AGS-1046081. We acknowledge David O’Neal of the Pittsburgh Supercomputing Center (PSC) for his assistance with the use of Tuning and Analysis Utilities (TAU) on a PSC Altix cluster early in the work. Computations were performed at the PSC and the National Institute for Computational Sciences (NICS), and the OU Supercomputing Center for Education and Research (OSCER).

REFERENCES

- Anderson, J. L., 2001: An ensemble adjustment Kalman filter for data assimilation. *Mon. Wea. Rev.*, **129**, 2884–2903.
- , 2003: A local least square framework for ensemble filtering. *Mon. Wea. Rev.*, **131**, 634–642.
- , and N. Collins, 2007: Scalable implementations of ensemble filter algorithms for data assimilation. *J. Atmos. Oceanic Technol.*, **24**, 1452–1463.
- Bishop, C. H., B. J. Etherton, and S. J. Majumdar, 2001: Adaptive sampling with the ensemble transform Kalman filter. Part I: Theoretical aspects. *Mon. Wea. Rev.*, **129**, 420–436.
- Brewster, K., M. Hu, M. Xue, and J. Gao, 2005: Efficient assimilation of radar data at high resolution for short-range numerical weather prediction. *Extended Abstracts, World Weather Research Programme Int. Symp. on Nowcasting and Very Short Range Forecasting*, Toulouse, France, Météo-France and Eumetsat, 3.06. [Available online at <http://www.meteo.fr/cic/wsn05/DVD/index.html>.]
- Buehner, M., P. L. Houtekamer, C. Charette, H. L. Mitchell, and B. He, 2010: Intercomparison of variational data assimilation and the ensemble Kalman filter for global deterministic NWP. Part II: One-month experiments with real observations. *Mon. Wea. Rev.*, **138**, 1567–1586.
- Burgers, G., P. J. van Leeuwen, and G. Evensen, 1998: Analysis scheme in the ensemble Kalman filter. *Mon. Wea. Rev.*, **126**, 1719–1724.
- Campbell, W. F., C. H. Bishop, and D. Hodyss, 2010: Vertical covariance localization for satellite radiances in ensemble Kalman filters. *Mon. Wea. Rev.*, **138**, 282–290.
- Caya, A., J. Sun, and C. Snyder, 2005: A comparison between the 4D-VAR and the ensemble Kalman filter techniques for radar data assimilation. *Mon. Wea. Rev.*, **133**, 3081–3094.
- Courtier, P., and O. Talagrand, 1987: Variational assimilation of meteorological observations with the adjoint equation. Part II: Numerical results. *Quart. J. Roy. Meteor. Soc.*, **113**, 1329–1347.
- Dong, J., M. Xue, and K. K. Droegemeier, 2011: The analysis and impact of simulated high-resolution surface observations in addition to radar data for convective storms with an ensemble Kalman filter. *Meteor. Atmos. Phys.*, **112**, 41–61.
- Dowell, D. C., F. Zhang, L. J. Wicker, C. Snyder, and N. A. Crook, 2004: Wind and temperature retrievals in the 17 May 1981 Arcadia, Oklahoma, supercell: Ensemble Kalman filter experiments. *Mon. Wea. Rev.*, **132**, 1982–2005.
- , L. J. Wicker, and C. Snyder, 2011: Ensemble Kalman filter assimilation of radar observations of the 8 May 2003 Oklahoma City supercell: Influence of reflectivity observations on storm-scale analysis. *Mon. Wea. Rev.*, **139**, 272–294.
- Evensen, G., 1994: Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics. *J. Geophys. Res.*, **99** (C5), 10 143–10 162.
- , 2003: The ensemble Kalman filter: Theoretical formulation and practical implementation. *Ocean Dyn.*, **53**, 343–367.
- , and P. J. van Leeuwen, 1996: Assimilation of Geosat altimeter data for the Agulhas Current using the ensemble Kalman filter with a quasigeostrophic model. *Mon. Wea. Rev.*, **124**, 85–96.
- Gao, J., M. Xue, Z. Wang, and K. K. Droegemeier, 1998: The initial condition and explicit prediction of convection using ARPS adjoint and other retrievals methods with WSR-88D data. Preprints, *12th Conf. on Numerical Weather Prediction*, Phoenix, AZ, Amer. Meteor. Soc., 176–178.
- Hamill, T. M., J. S. Whitaker, M. Fiorino, and S. G. Benjamin, 2011: Global ensemble predictions of 2009’s tropical cyclones initialized with an ensemble Kalman filter. *Mon. Wea. Rev.*, **139**, 668–688.
- Houtekamer, P. L., and H. L. Mitchell, 1998: Data assimilation using an ensemble Kalman filter technique. *Mon. Wea. Rev.*, **126**, 796–811.
- Jung, Y., M. Xue, G. Zhang, and J. Straka, 2008: Assimilation of simulated polarimetric radar data for a convective storm using ensemble Kalman filter. Part II: Impact of polarimetric data on storm analysis. *Mon. Wea. Rev.*, **136**, 2246–2260.

- , —, and M. Tong, 2012: Ensemble Kalman filter analyses of the 29–30 May 2004 Oklahoma tornadic thunderstorm using one- and two-moment bulk microphysics schemes, with verification against polarimetric data. *Mon. Wea. Rev.*, **140**, 1457–1475.
- Keppenne, C. L., and M. M. Rienecker, 2002: Initial testing of a massively parallel ensemble Kalman filter with the Poseidon isopycnal ocean general circulation model. *Mon. Wea. Rev.*, **130**, 2951–2965.
- Le Dimet, F. X., and O. Talagrand, 1986: Variational algorithms for analysis and assimilation of meteorological observations: Theoretical aspects. *Tellus*, **38A**, 97–110.
- Li, X., and M. Parashar, 2007: Hybrid runtime management of space-time heterogeneity for parallel structured adaptive applications. *IEEE Trans. Parallel Distrib. Syst.*, **18**, 1202–1214.
- Liu, H., and M. Xue, 2006: Retrieval of moisture from slant-path water vapor observations of a hypothetical GPS network using a three-dimensional variational scheme with anisotropic background error. *Mon. Wea. Rev.*, **134**, 933–949.
- Michalakes, J., J. Dudhia, D. Gill, T. Henderson, J. Klemp, W. Skamarock, and W. Wang, 2004: The Weather Research and Forecast model: Software architecture and performance. *Proc. 11th Workshop on the Use of High Performance Computing in Meteorology*, Reading, United Kingdom, ECMWF, 156–168.
- Milbrandt, J. A., and M. K. Yau, 2005: A multimoment bulk microphysics parameterization. Part I: Analysis of the role of the spectral shape parameter. *J. Atmos. Sci.*, **62**, 3051–3064.
- Parashar, M., X. Li, and S. Chandra, 2010: *Advanced Computational Infrastructures for Parallel and Distributed Adaptive Applications*. John Wiley & Sons, 518 pp.
- Ray, P. S., B. Johnson, K. W. Johnson, J. S. Bradberry, J. J. Stephens, K. K. Wagner, R. B. Wilhelmson, and J. B. Klemp, 1981: The morphology of several tornadic storms on 20 May 1977. *J. Atmos. Sci.*, **38**, 1643–1663.
- Sakov, P., G. Evensen, and L. Bertino, 2010: Asynchronous data assimilation with the EnKF. *Tellus*, **62**, 24–29.
- Sathye, A., M. Xue, G. Bassett, and K. K. Droegemeier, 1997: Parallel weather modeling with the Advanced Regional Prediction System. *Parallel Comput.*, **23**, 2243–2256.
- Snook, N., M. Xue, and J. Jung, 2011: Analysis of a tornadic mesoscale convective vortex based on ensemble Kalman filter assimilation of CASA X-band and WSR-88D radar data. *Mon. Wea. Rev.*, **139**, 3446–3468.
- Snyder, C., and F. Zhang, 2003: Assimilation of simulated Doppler radar observations with an ensemble Kalman filter. *Mon. Wea. Rev.*, **131**, 1663–1677.
- Sun, J., and N. A. Crook, 1997: Dynamical and microphysical retrieval from Doppler radar observations using a cloud model and its adjoint. Part I: Model development and simulated data experiments. *J. Atmos. Sci.*, **54**, 1642–1661.
- Szunyogh, I., E. J. Kostelich, G. Gyarmati, E. Kalnay, B. R. Hunt, E. Ott, and E. Satterfield, 2008: A local ensemble transform Kalman filter data assimilation system for the NCEP global model. *Tellus*, **60A**, 113–130.
- Tippett, M. K., J. L. Anderson, C. H. Bishop, T. M. Hamill, and J. S. Whitaker, 2003: Ensemble square root filters. *Mon. Wea. Rev.*, **131**, 1485–1490.
- Tong, M., and M. Xue, 2005: Ensemble Kalman filter assimilation of Doppler radar data with a compressible non-hydrostatic model: OSS experiments. *Mon. Wea. Rev.*, **133**, 1789–1807.
- , and —, 2008: Simultaneous estimation of microphysical parameters and atmospheric state with simulated radar data and ensemble square root Kalman filter. Part II: Parameter estimation experiments. *Mon. Wea. Rev.*, **136**, 1649–1668.
- Wang, S., M. Xue, and J. Min, 2013: A four-dimensional asynchronous ensemble square-root filter (4DEnSRF) algorithm and tests with simulated radar data. *Quart. J. Roy. Meteor. Soc.*, **139**, 805–819.
- Whitaker, J. S., and T. M. Hamill, 2002: Ensemble data assimilation without perturbed observations. *Mon. Wea. Rev.*, **130**, 1913–1924.
- Wu, B., J. Verlinde, and J. Sun, 2000: Dynamical and microphysical retrievals from Doppler radar observations of a deep convective cloud. *J. Atmos. Sci.*, **57**, 262–283.
- Xue, M., and Coauthors, 1996: The 1996 CAPS spring operational forecasting period: Realtime storm-scale NWP, Part II: Operational summary and examples. Preprints, *11th Conf. Numerical Weather Prediction*, Norfolk, VA, Amer. Meteor. Soc., 297–300.
- , K. K. Droegemeier, and V. Wong, 2000: The Advanced Regional Prediction System (ARPS)—A multiscale non-hydrostatic atmospheric simulation and prediction tool. Part I: Model dynamics and verification. *Meteor. Atmos. Phys.*, **75**, 161–193.
- , and Coauthors, 2001: The Advanced Regional Prediction System (ARPS)—A multi-scale nonhydrostatic atmospheric simulation and prediction tool. Part II: Model physics and applications. *Meteor. Atmos. Phys.*, **76**, 143–165.
- , D.-H. Wang, J. Gao, K. Brewster, and K. K. Droegemeier, 2003: The Advanced Regional Prediction System (ARPS), storm-scale numerical weather prediction and data assimilation. *Meteor. Atmos. Phys.*, **82**, 139–170.
- , M. Tong, and K. K. Droegemeier, 2006: An OSSE framework based on the ensemble square root Kalman filter for evaluating the impact of data from radar networks on thunderstorm analysis and forecasting. *J. Atmos. Oceanic Technol.*, **23**, 46–66.
- , K. K. Droegemeier, and D. Weber, 2007: Numerical prediction of high-impact local weather: A driver for petascale computing. *Petascale Computing: Algorithms and Applications*, Taylor & Francis Group, LLC, 103–124.
- , M. Tong, and G. Zhang, 2009: Simultaneous state estimation and attenuation correction for thunderstorms with radar data using an ensemble Kalman filter: Tests with simulated data. *Quart. J. Roy. Meteor. Soc.*, **135**, 1409–1423.
- , Y. Jung, and G. Zhang, 2010: State estimation of convective storms with a two-moment microphysics scheme and an ensemble Kalman filter: Experiments with simulated radar data. *Quart. J. Roy. Meteor. Soc.*, **136**, 685–700.
- , and Coauthors, 2011: Realtime convection-permitting ensemble and convection-resolving deterministic forecasts of CAPS for the Hazardous Weather Testbed 2010 Spring Experiment. *Extended Abstracts, 24th Conf. on Weather Forecasting/20th Conf. Numerical Weather Prediction*, Seattle, WA, Amer. Meteor. Soc., 9A.2. [Available online at https://ams.confex.com/ams/91Annual/webprogram/Manuscript/Paper183227/Xue_CAPS_2011_SpringExperiment_24thWAF20thNWP_ExtendedAbstract.pdf.]
- Zhang, S., M. J. Harrison, A. T. Wittenberg, A. Rosati, J. L. Anderson, and V. Balaji, 2005: Initialization of an ENSO forecast system using a parallelized ensemble filter. *Mon. Wea. Rev.*, **133**, 3176–3201.