Dear Editor,

Herewith we submit the revised manuscript gmd-2016-208, " Accelerating volcanic ash data assimilation using a mask-state algorithm based on ensemble Kalman filter: a case study with the LOTOS-EUROS model (version 1.10) ".

First we would like to thank you and three anonymous reviewers, and really appreciate the detailed comments and suggestions. We have carefully considered the comments and made changes accordingly in the revised paper. We believe the new version has been improved a lot compared to the previous version. In the new version, with the agreement of all authors, we have changed the author's order to Guangliang Fu, Hai Xiang Lin, Arnold Heemink, Sha Lu, Arjo Segers, Nils van Velzen, Tongchao Lu, Shiming Xu.

In the following we will give our answers and reactions to the comments. To make the changes easier to identify, we have numbered them.

Best regards,
Guangliang Fu
on behalf all co-authors

(The revised manuscript is in the latter part of this pdf.
Changes for Reviewer#1: "purple"; Reviewer#2: "red"; Reviewer#3: "blue".)

1

Reply to Reviewer #1:

Reply to **Specific comments**:

1. *The variable notation N_{ash} should be changed to n_{ash} as the small n presents number of grid points (line 4.30).*

   Response:

   I agree. All N_{ash} have been changed to n_{ash} in the revision.

2. *Equation (15) indicates whether we would have speedup by the MS method (line 4.5). Could it be extended to explicitly specify the expected amount of speed up (with some assumptions if necessary)?*

   Response:

   Thanks for the comment. As suggested, it is extended in **line(s) 9.19–9.23**:

   The relation $n_{ash} < \frac{N-1}{N} n$ indicates whether we would have speedup by the MS method, actually it can be extended to Eq. (15),

   $$S_{\mathrm{ms}} = \frac{\mathrm{O}(n\,N^2)}{\mathrm{O}(n\,N + n_{ash}\,N^2)} = \mathrm{O}(\frac{n}{n_{ash}}), \tag{15}$$

   which explicitly specifies the expected amount of speedup (S_{ms}) of $\mathbf{A}^a = \mathbf{A}^f\,\mathbf{X}$ by the MS algorithm. In this case study, N is taken at 100 and $\frac{n_{ash}}{n} \approx \frac{1}{3}$, so S_{ms} is approximately 3.0.

3. *Some explanations are necessary about the involved operations in the others part for the clarification and justification of the less (by 0.03h) computational time (Table-1).*

   Response:

   I agree. The explanations are added in **line(s) 10.19–10.23**:

   Note that it can also be observed that the computational time for the "other" parts in Table 1 (such as operations for setting environmental variables, starting and finalizing DA algorithms, as mentioned in Sect. 3.1) is slightly reduced by the MS method (i.e., 0.03 h in this case). This is because in the conventional EnKF, the ensemble mean state $\bar{\mathbf{A}}^f$ is calculated in the "other" parts as an output to finalize the DA algorithms, while in MS-EnKF, the calculations of $\bar{\mathbf{A}}^f$ are needed and directly involved in the "Analysis" part.

4. *The computational overhead involved to transform the full matrix to a small one and vice versa should be explicitly discussed. The effect in the computation time should be also discussed.*

   Response:

   As suggested, the computational overhead involved to transform the full matrix to a small one and vice versa are explicitly discussed in **line(s) 8.12–8.14**:

   Construct masked ensemble state matrix $\tilde{\mathbf{A}}^f$: Using the mask array $\mathbf{z}_{n_{ash} \times 1}$ obtained from step (ii), $\tilde{\mathbf{A}}^f_{n_{ash} \times N}$ can be constructed column by column according to Eq. (13), and the computational cost (overhead) for this step is $\mathrm{O}(n_{ash}\,N)$.

   $$\tilde{\mathbf{A}}^f(1 : n_{ash}, 1 : N) = \mathbf{A}^f(\mathbf{z}(1 : n_{ash}), 1 : N), \tag{13}$$

and in **line(s) 8.17–8.20**:

> **Construct analyzed ensemble state matrix $\mathbf{A}^a$**: With the computed $\tilde{\mathbf{A}}^a$ from step (iv) and the mask array $\mathbf{z}$ from step (ii), the final analyzed ensemble state matrix $\mathbf{A}^a_{n \times N}$ can be constructed based on Eq. (14). The computational cost (overhead) for this step is $\mathrm{O}(n\,N)$.
>
> $$\mathbf{A}^a(\mathbf{z}(1:n_{ash}),1:N) = \tilde{\mathbf{A}}^{\mathbf{a}}(1:n_{ash},1:N), \tag{14}$$

The effect in the computation time is discussed in **line(s) 9.5–9.10**:

> Now we analyze MS's computational cost, which can be summed as $\mathrm{O}(n\,N)+\mathrm{O}(n)+\mathrm{O}(n_{ash}\,N)+\mathrm{O}(n_{ash}\,N^2)+\mathrm{O}(n\,N)$, i.e., $\mathrm{O}(n\,N+n_{ash}\,N^2)$. Thus, the computational overhead involved to transform the full matrix to a small one (i.e., $\mathrm{O}(n_{ash}\,N)$ for procedure (iii)) has little effect in the total computation cost of MS (i.e., $\mathrm{O}(n\,N + n_{ash}\,N^2)$). However, the computational overhead of transforming the small matrix to the full one (i.e., $\mathrm{O}(n\,N)$ for procedure (v)) does contribute a part, which cannot be ignored, to the total MS's computational cost.

5. *Line 3.20 : It is mentioned that 72% of the total runtime is spent on the analysis step. This information can be further extended (preferably with mathematical formulation) to predict the maximum computational speedup (e.g. 3.57 for the case in this paper). c.f. Amdahl's law (Amdahl, Gene M. (1967). "Validity of the Single Processor Approach to Achieving Large-Scale Computing Capabilities". AFIPS Conference Proceedings (30): 483-485.)*

6. *Related to the above point, a further discussion about the discrepancy in the actual speedup (e.g. 2.24 for the case in this paper) from the predicted speedup could be included.*

Response:

I agree. An extended discussion is added in **line(s) 9.24–9.32**:

> According to Amdahl's law (Amdahl, 1967), the total computational speedup (S_{total}) by MS can be predicted by Eq. (16),
>
> $$S_{\mathrm{total}} = \frac{1}{(1 - p_{\mathrm{ms}}) + \frac{p_{\mathrm{ms}}}{S_{\mathrm{ms}}}}, \tag{16}$$
>
> where p_{ms} is the proportion of the computational cost of $\mathbf{A}^a = \mathbf{A}^f\,\mathbf{X}$ in the overall DA computations. It has been evaluated that the computational cost of $\mathbf{A}^a = \mathbf{A}^f\,\mathbf{X}$ dominates the analysis step (see Fig. 2(b)), thus the proportion of the computational cost of $\mathbf{A}^a = \mathbf{A}^f\,\mathbf{X}$ approximates the proportion of the analysis step in the total DA computations (i.e., $p_{\mathrm{ms}} \approx 72\,\%$ in this case, as described in Sect. 3.1). Therefore, based on Eq. (16), the maximum ("ideal") computational speedup can be predicted to be $\frac{1}{1-p_{\mathrm{ms}}}$ (i.e., ≈ 3.57 for this case study) when S_{ms} approximates infinity. However, this is not the actual speedup because S_{ms} is in fact specified by Eq. (15). Based on discussions above, S_{total} can be therefore estimated by Eq. (15) at ≈ 2.0 in this case.

7. *It is understood from equation (15) the total speedup is quite low, is there any statistical data regarding the value of n_{ash}/n to know the probability of the applicability of the proposed method and justify the weight of the work?*

Response:

Thanks for the suggestion. The related discussion is added in **line(s) 14.17–14.23**:

As stated in Eq. (15), the speedup of the MS method is approximately the inverse of $\frac{n_{ash}}{n}$. So far there is no statistical data on the value of $\frac{n_{ash}}{n}$. Considering the problem of volcanic ash transport, there is one emission point (at the volcano), all the ashes in atmospheres are transported by the directional wind drive from the same source point. Thus volcanic ash cloud is actually transported in a shape of a plume, which in general doesn't cover the full but only a small part of the 3-D domain. At the start phase of a volcanic ash eruption, $\frac{n_{ash}}{n}$ is much smaller than 1.0 (started from 0). During transport over a long time (one and a half months for this case study), $\frac{n_{ash}}{n}$ increases to approximately $\frac{1}{3}$. Therefore, the speedup of MS on volcanic ash DA will be significant.

8. *The paper investigates actually a single method (MS) to accelerate volcanic ash data assimilation. Other methods are only referred (e.g. parallelization of the analysis step) or contrasted (e.g. localization) and are not really investigated in this paper. So, the plural form (e.g. strategies) should be omitted.*

Response:

I agree. All "strategies" in the plural form have been changed to "a strategy" in the revision.

Reply to Reviewer #2:

Reply to comments:

1. *The most costly step identified by the authors is a simple matrix multiplication,*
   $\mathbf{A}^a = \mathbf{A}^f\,\mathbf{X}$. *Such a simple step should be very easily distributed over the available CPU's and should be very fast, so in fact I do not even understand the "problem" that the authors want to solve.*

   Response:
   The problem is "we aim to accelerate the runtime time to less than the simulation time", which is important for the assimilation in an operational sense. It is described in **line(s) 5.16–5.19**:

   The evaluation result of the conventional EnKF is shown in Table 1 (the middle column). It can be seen that the total computational time (4.36 h) is relatively large compared to the simulation window (3.0 h, i.e., from 9:00 - 12:00 UTC, 18 May, 2010), which is too much in an operational sense. Therefore, in this study, we aim to accelerate the computation to within an acceptable runtime (i.e., requires less runtime than the time period of the DA application).

   We agree that the time-consuming part $\mathbf{A^a} = \mathbf{A^f X}$ is not a complicated matrix multiplication. However, we have employed the fastest available CPU on Cartesius (Netherlands' supercomputer. Each node is configured with 2×12-core 2.6 GHz Intel Xeon E5-2690 v3 (Haswell) CPUs and with memory 64 GB), the total computational time (4.36 h) is still not acceptable, compared to a simulation time of 3 h. This is mainly because $\mathbf{A}^a = \mathbf{A}^f\,\mathbf{X}$ takes large amounts of time for the analysis step of 3.14 h. Thus, that motivated us to develop the Mask-State algorithm and the result shows it is efficient for our case. Therefore, our problem can be briefly understood as "a further acceleration of $\mathbf{A^a} = \mathbf{A^f X}$" based on model dynamics.

   In order to speed up the computation, we can do two things: (1) reduce the amount of operations in $\mathbf{A}^f\,\mathbf{X}$; (2) using parallel computation. For the application at hand, the Mask-state algorithm strongly reduces the total amount of computation; parallel processing can be applied after that.

2. *This topic by itself is very technical, and arguably does not belong in the "geoscientific" journal GMD. The results may justify publication if the approach would be applicable to a large class of assimilation problems, but this is not the case as mentioned by the authors themselves. Using masks is only applicable for a very limited set of problems, typically single point source releases of short-lived species.*

   Response:
   Using Mask-State (MS) is applicable for many problems, where the domain is not fully polluted by the species. It is certainly not limited to "typically single point source releases of short-lived species". It doesn't matter what the emission looks like and whether the releases are "short-lived species" or "long-lived species". Given an assimilation problem, the only restriction for MS to gain an acceleration is whether the whole domain is fully polluted or partly polluted.

   MS is therefore suitable for many assimilation problems, such as all the volcanic-related ash/gas assimilation, e.g., assimilation of satellite data/LIDAR data/in situ data; (sand/desert) dust storm related assimilation; tornado-related assimilation; assimilation of exploding nuclear plants or factories; chemicals or oils leaking on seas; global (forecast) fire assimilation; assimilation of environmental pollutant transport, e.g., severe smog. In addition, for data assimilation applications (e.g., Ozone, SO_2) where pollutants spread over the whole domain, usually the focus is only on the high concentrations and a threshold can be set to ignore the very low values without losing the necessary

assimilation accuracy. In this case, MS can also lead to a potential acceleration since many very low concentrations can be explicitly truncated to be zeros.

For all the above applications, MS can accelerate $\mathbf{A}^a = \mathbf{A}^f \mathbf{X}$. As a clarification, we have clearly mentioned this in the new version in **line(s) 13.26–14.3**:

Using MS is also applicable for many other DA problems, where the domain is not fully polluted by the species. It does not matter what the emission looks like and whether the releases are short- or long-lived species. Given an assimilation problem, the only restriction for MS to gain an acceleration is whether the whole domain is fully polluted or partly polluted. The assimilation problems where MS can achieve the acceleration effect on the computations of $\mathbf{A}^a = \mathbf{A}^f \mathbf{X}$ include all the volcanic-related ash/gas assimilation, e.g., assimilation of satellite data/LIDAR data/in situ data; (sand/desert) dust storm related assimilation; tornado-related assimilation; assimilation of exploding nuclear plants or factories; chemicals or oils leaking on seas; global (forecast) fire assimilation; assimilation of environmental pollutant transport, e.g., severe smog. In addition, for DA applications (e.g., Ozone, SO_2) where pollutants spread over the whole domain, usually the focus is only on the high concentrations and a threshold can be set to ignore the very low values without losing the necessary assimilation accuracy. In this case, MS can also lead to a potential acceleration since many very low concentrations can be explicitly truncated to be zeros.

This study is submited as a "Development and technical paper" to GMD, focusing on the development of MS to accelerate a data assimilation application not only for the volcanic ash forecast problem which is used as a case study, but also for many other assimilation problems. Therefore, we believe it fits the scope of GMD.

3. *The problem is very idealised, with only a few observations (m=2). In this case the whole inversion problem in "observation space" is computationally very fast, which normally is not the case.*

Response:
Fu et al. (2016) described the methodology for an aircraft-based volcanic ash data assimilation problem, where (for one aircraft) at each assimilation time, two in situ measurements of PM_{10} and $PM_{2.5}$ are assimilated (thus $m=2$ here). The problem should not be considered as an idealised case, but an example of stationary data assimilation (although aircraft is not exactly stationary, but the measurement number at each assimilation time is steady and stationary).

We admit for $m=2$, the whole inversion problem in "observation space" is computationally fast. This is shown in Fig. 2(b), where at one analysis step, the computational cost of the inversion problem ($\mathbf{X_4} = \mathbf{X_3}^{-1}$) is O($m^3$). However, for stationary data assimilation problems where mostly $m \ll n$ (n is the state number, in this case $n \sim 10^6$), the inversion problem is also fast compared to the part $\mathbf{A}^a = \mathbf{A}^f \mathbf{X}$ (O($n N^2$), N=100 is the ensemble size). Thus, the case $m=2$ can be taken as an example of stationary assimilation. Actually, this type of assimilation is a common assimilation type. For example, Barbu et al. (2009) assimilated measurements of SO_2 and SO_4 from the EMEP database. The assimilation set contains 17 sites for SO_2 and 27 for SO_4, thus $m=44$ ($m \ll n$).

Recently, a large number of national weather services have implemented ceilometer networks, mainly for monitoring the dispersion of volcanic ash clouds (Wiegner et al., 2014). These data set will be (and in part are already) available in near real time and will provide information about the (horizontal and) vertical distribution (with some restrictions due to cloud cover). Thus, they could be very promising candidates for stationary data assimilation for volcanic ash plumes. Under these stationary measurement setup ($m \sim 150$), based on Fig. 2(b), the inversion problem is still not expensive (O(m^3)) compared to $\mathbf{A}^a = \mathbf{A}^f \mathbf{X}$ (O($n N^2$)) at one analysis step.

Therefore, for stationary/near-stationary data assimilation, it is usually the case that the inversion problem in "observation space" is computationally fast.

We also admit the problem discussed in the study is not the case where O(m)=O(n), e.g., for passive satellite data assimilation, where the inversion problem in the analysis step is the most time-consuming part (O(m^3)=O(n^3)), but this is not the case considered in this study for stationary data assimilation. We have added some related discussion in **line(s) 17.2–17.7**:

6

However, in some applications when many measurements are assimilated (e.g., satellite-based or seismic-based data), and the number of measurements is of the same order as the number of state variables, the most time-consuming part will be the SVD. In these cases, the contributions of MS will be limited. The reduction of the total computing time using MS therefore is less significant, an effective acceleration algorithm for the analysis step must be used and should consider the computationally-expensive SVD in the first place.

Reply to Reviewer #3:

Reply to **general comments**:

1. *More numerically sophisticated areas like the modeling community will find this standard but it is likely to have an impact on the DA community so I recommend publication after the authors carefully make the case for the Mask State scheme relative to more standard sparse matrix methods not just comparison to full storage dense matrices.*

   Response:

   We agree that the previous only comparison of the mask-state algorithm (MS) to the cases of full storage dense matrices was not enough. The comparison of MS to more standard sparse matrix methods (e.g., CSR-based sparse-dense matrix multiplication (SDMM)) is necessary, because MS can also be taken as one specific sparse matrix method, which typically works for ensemble-based data assimilation (DA) applications.

   We have included a comparison in the revision where the case of MS is compared with standard sparse matrix methods. The detailed comparison and explanations are presented in the new Sect. 5 (two new tables (Table 3 to 4) are added to show the result), in **line(s) 11.16–13.22**, as shown in the following.

5 Comparison between MS and standard sparse matrix methods

5.1 Issues related to the generation of CSR-based arrays

According to Sect. 4, MS has proven to be capable of solving the computational issue of $\mathbf{A}^a = \mathbf{A}^f \mathbf{X}$. Motivated from model's characteristics, MS was proposed from an application's perspective and achieved a good result by managing the irregular sparsity in our complicated volcanic ash DA system. The main reason why MS is efficient is that the sparsity of $\mathbf{A}^f$ can be well utilized by MS. In Sect. 4, we only performed the comparison between MS and the case of full storage dense matrices. However, the problem abstracted here ($\mathbf{A}^f \mathbf{X}$) is actually a sparse-dense matrix multiplication (SDMM) problem, since $\mathbf{A}^f$ is sparse and $\mathbf{X}$ is dense (see Eq. (12) for $\mathbf{X}$). Thus, one may wonder what the result would be, if the comparison of MS is made to more standard sparse matrix methods, such as Compressed Sparse Row (CSR)-based methods (Saad, 2003; Bank and Douglas, 1993), which are commonly used for sparse matrix vector/matrix multiplication.

Before we make the comparison, we need to first address the intrinsic problem when considering standard sparse matrix methods in EnKF for $\mathbf{A}^a = \mathbf{A}^f \mathbf{X}$. The issue is that it is not possible to directly generate a sparse storage format (e.g., CSR) of $\mathbf{A}^f$ without first generating the full matrix $\mathbf{A}^f$. This is mainly because $\mathbf{A}^f$ comes from the model-driven ensemble forecast step, where each $\mathbf{A}^f$ column corresponds to one member of the ensemble. During model forecast, we know there are indeed no-ash grids. However, it is not sure where the plume exactly is after one forecasting time step. This is highly dependent on the weather conditions and the model processes (e.g, advection and diffusion for horizontal grids, sedimentation and deposition for vertical grids). Thus, a fixed and wide domain is usually needed by the model to avoid complications, resulting in the generation of the full storage of $\mathbf{A}^f$ (to be used in $\mathbf{A}^a = \mathbf{A}^f \mathbf{X}$). Therefore, if we want to implement a CSR storage format for the sparse matrix $\mathbf{A}^f$, we must first generate the full storage $\mathbf{A}^f$ from the ensemble forecast step, and then we generate the three CSR arrays based on $\mathbf{A}^f$.

Generating CSR arrays is usually much more expensive (computationally) than a single sparse matrix-vector multiplication (SpMV). Thus, if we generate CSR arrays for only performing one-time SpMV, it would be meaningless from HPC's point of view. Fortunately, this is not the case

for $\mathbf{A}^a = \mathbf{A}^f\,\mathbf{X}$ (i.e., SDMM), which can actually be considered as N-times SpMV. (Here, $\mathbf{X}$ has N columns, and one SpMV means the multiplication of $\mathbf{A}^f$ with one column of $\mathbf{X}$.) Thus, CSR-based SDMM might also be a candidate in reducing the computation time of $\mathbf{A}^a = \mathbf{A}^f\,\mathbf{X}$. It remains interesting to compare the performance of CSR-based SDMM and MS, in dealing with $\mathbf{A}^a = \mathbf{A}^f\,\mathbf{X}$ for our study case.

5.2   Result of CSR-based SDMM

To implement CSR-based SDMM for $\mathbf{A}^a = \mathbf{A}^f\,\mathbf{X}$, the three CSR arrays for $\mathbf{A}^f$ (denoted $\boldsymbol{val}$, $\boldsymbol{col\_idx}$, $\boldsymbol{row\_ptr}$ in this study) need to be first generated. The array $\boldsymbol{val}$ of size $n_{\boldsymbol{val}}$ stores non-zero values of $\mathbf{A}^f$, where $n_{\boldsymbol{val}} \approx n_{ash}\,N$. (In this study case, $n = 3.888 \times 10^6$, $n_{ash} \approx \frac{1}{3}\,n$, $N = 100$.) The array $\boldsymbol{col\_idx}$ of the same size $n_{\boldsymbol{val}}$ stores column index of the non-zeros. The array $\boldsymbol{row\_ptr}$ saves the start and the end pointers of the non-zeros of the rows in $\mathbf{A}^f$. The size of $\boldsymbol{row\_ptr}$ is $n + 1$.

After the above three CSR arrays are generated, CSR-based SpMV can be performed for multiplying $\mathbf{A}^f$ with a column vector $\boldsymbol{v}$ in $\mathbf{X}$ (see Algorithm 1 in Fig. 4(a)). After that, Algorithm 2 (Fig. 4(b)) can be implemented by looping Algorithm 1 for N times to obtain $\mathbf{A}^a = \mathbf{A}^f\,\mathbf{X}$. The experimental result of CSR-based SDMM is shown in Table 4, where all the environmental conditions (such as the DA system, the programming environment) are the same as the case of MS. This gives a fair comparison between CSR-based SDMM and MS. In addition, for a pure algorithmic comparison with the serial MS, here the CSR-based SDMM is also performed in a serial case.

From Table 4, we can first confirm that the computational time (i.e., 0.0407 h) for the generation of the three CSR-based arrays ($\boldsymbol{val}$, $\boldsymbol{col\_idx}$, $\boldsymbol{row\_ptr}$ to represent the sparse matrix $\mathbf{A}^f$) indeed takes more time than the computational time of one CSR-based SpMV (i.e., 0.0117 h). Thus, there is little value to perform sub-step (i) (see Table 4) if only one SpMV (i.e., sub-step (ii)) is needed. However, to get all N (i.e., 100) columns of $\mathbf{A}^a$, the sub-step (ii) is looped for N times, resulting in an ignorable impact of sub-step (i) on the total computational time (i.e., 1.21 h) of CSR-based SDMM.

The result of CSR-based SDMM also shows that the standard sparse matrix methods can reduce the computational time of $\mathbf{A}^a = \mathbf{A}^f\,\mathbf{X}$, by comparing with the conventional way in Table 3. However, it can also be observed that the computational time of CSR-based SDMM is larger than MS (i.e., 1.21 h versus 0.87 h in Table 3). Thus, although application of sparse matrix multiplication methods is positive, it is still slower than MS on our problem.

5.3   Comparison between CSR-based SDMM and MS

In the CSR-based SDMM, only non-zero elements in $\mathbf{A}^f$ participate in the multiplication between $\mathbf{A}^f$ and $\mathbf{X}$, thus redundant computation (related to zero elements in $\mathbf{A}^f$) is avoided. So the computation time of $\mathbf{A}^f\,\mathbf{X}$ is reduced with CSR-based SDMM. In the following, we analyze the performance difference between CSR-based SDMM and MS.

Firstly, from the programming's perspective, in CSR-based SDMM, the loop number for the rows of $\mathbf{A}^f$ is from 1 to n (see Fig. 4(a)), while the corresponding loop number in MS is from 1 to n_{ash} (see the step (iv) of MS in Sect. 4.2, $n_{ash} \approx \frac{1}{3}n$). Although only non-zero elements are used in the multiplication in CSR-based SDMM, the length of the outer loop is still n (much larger than n_{ash}) which is the essential reason that MS is faster than CSR-based SDMM. Note that as discussed in Sect. 4.1, there are many zero rows in $\mathbf{A}^f$, thus CSR-based SDMM actually does nothing when it comes to a zero row, but still needs to execute the loop. Within each loop number, it has to check the information from $\boldsymbol{row\_ptr}$ (size $n + 1$), where the value corresponding to a zero row is usually set to be the value in $\boldsymbol{row\_ptr}$ corresponding to the first following non-zero row.

Secondly, with respect to the algorithm, CSR-based SDMM utilizes the sparsity of $\mathbf{A}^f$ by its generation of three CSR arrays. While, MS not only utilizes the sparsity information of the sparse matrix $\mathbf{A}^f$, but also utilizes the consistency of ensemble forecasts, that is, ensemble forecasted states are not consistent in values but usually consistent in non-zero locations. This is a typical property in ensemble-based DA, resulting in that N ensemble plumes are different in concentration values,

but have similar transport directions/shapes (see Sect. 4.2). Thus, most of the zero elements in $\mathbf{A}^f$ are actually in zero rows of $\mathbf{A}^f$ for an EnKF application, which leads to a small number of non-zero rows (n_{ash}) compared to the full number of rows (n) of $\mathbf{A}^f$. Therefore, the way only considering n_{ash} rows in $\tilde{\mathbf{A}}^f_{n_{ash} \times N} \mathbf{X}_{N \times N}$ (see the step (iv) of MS in Sect. 4.2) is more advantageous for an EnKF application than that considering all n rows in CSR-based SDMM. Based on the above analysis, MS can be considered as a specific sparse matrix method, which typically works for ensemble-based DA applications.

It is useful to apply standard sparse matrix methods (e.g., CSR-based SDMM) for our assimilation application. The accelerated analysis step by CSR-based SDMM (1.22 h, see Table 3) also reduces the total computational time (i.e., 2.29 h, see Table 1 for the computational time of initialization, forecast and others) to an acceptable level (i.e., less than 3 h for our case study). In practice, due to the better performance of MS than CSR-based SDMM, we will use MS as a better choice for assimilation applications. In addition, we do not only intend to present MS, but also intend to reveal which part is the most time-consuming part for plume-type assimilation of in situ observations.

2. *If the authors were to attack or even attempt to tackle the harder problem of programming for parallelism with changing sparsity this would be also be a much stronger and impactful paper.*

Response:

In the revision, we have added Sect. 6.3 to discuss on the parallelization issues. In this paper, for the current usage, we keep the possibility of parallelization open, becuase a serial MS has been shown to be efficient already. Therefore, although (from application's aspect) further acceleration by parallelization is not urgent, we fully agree that a detailed discussion on parallelization in MS or CSR-based SDMM is useful and will have a good added value to the presentation of the serial MS. The discussion is in **line(s) 15.19–16.8** in the paper as follows,

6.3 MS and parallelization

Motivated from the model's physics, the implementation MS currently is for the serial case. This implementation has reduced the computation time to an acceptable time (i.e., simulation time is less than the period of forecast in real world time). It is however interesting to discuss the potential of parallelization of the dense-dense matrix multiplication ($\tilde{\mathbf{A}}^a_{n_{ash} \times N} = \tilde{\mathbf{A}}^f_{n_{ash} \times N} \mathbf{X}_{N \times N}$) in the step (iv) of the algorithm (see Sect. 4.2 and Table 2). The related matrix multiplication can be easily parallelized on multiple processors. Optimization and evaluation on the parallelized MS will be considered in future. For the current case study, the computational time (3.13 h, see Table 3) for an "ideal" reduction by parallelization of MS is not much larger than the acceleration (already) gained by MS (2.26 h, subtraction between 3.13 h and 0.87 h, see Table 3). Therefore, from application's perspective, further acceleration by parallelization is not required.

Alternatively, one may also consider to (1) directly parallelize the expensive matrix multiplication of $\mathbf{A}^a_{n \times N} = \mathbf{A}^f_{n \times N} \mathbf{X}_{N \times N}$, without first performing MS; or (2) implement CSR-based SDMM (see Sect. 5) with parallelization. Both are possible alternative approaches to accelerate the expensive matrix multiplication. The first approach can be implemented by user own designed parallelization, or by utilizing scaLAPACK (https://www.netlib.org/scalapack/, where the main function is "pdgemm"). The second approach can be realized by using some general parallel sparse-dense matrix multiplication methods (e.g., sending each column of X and three CSR arrays of $\mathbf{A}^f$ to each processor to calculate each column of $\mathbf{A}^a$) or using a good parallel algebra library like PeTSC (https://www.mcs.anl.gov/petsc/) which allows users to specify own orderings and comes with machine optimized parallel matrix-matrix multiplication operations. However, given the fact that MS can also be parallelized using similar ways or same libraries, thus it is fair to not consider parallelization for all cases (i.e., using MS, not using MS, using CSR-based SDMM). Actually, the parallelization in MS could be performed much easier than other approaches in dealing with $\mathbf{A}^a = \mathbf{A}^f \mathbf{X}$, because the dense-dense matrix multiplication (parallelization in the step (iv) of MS) is easier to parallelize than the sparse-dense matrix multiplication (direct parallelization for

$\mathbf{A}^a = \mathbf{A}^f \, \mathbf{X}$ or parallelized CSR-based SDMM).

In this paper, for the current usage, we keep the possibility of parallelization open, because a serial MS has been efficient already.

In addition, in the ash forecast step, due to the changing sparsity in $\mathbf{A}^f$, a fixed domain-decomposed parallelization to compute $\mathbf{A}^f$ usually would not work efficiently. In the future if further parallelization is required to deal with large domain and more complicated forecast step, some advanced approach such as adaptive domain-decomposed parallelization (Lin et al., 1998) might be employed. We have also discussed on this issue in **line(s) 11.9–11.15**:

As for domain-decomposed parallelization (Segers, 2002), it is not efficient for our application. This is because volcanic ash is special in the sense that the model is only doing computations in a small part of the domain (i.e., there is no data in a rather large part of domain), and this active part is continuously changing. Thus, a fixed domain decomposition is not very useful here because of the changing plume position. In this sense, some advanced approach such as adaptive domain-decomposed parallelization (Lin et al., 1998) should be adopted to achieve additional acceleration to the volcanic ash forecast stage. This is an interesting subject for future application, when a more complicated model is employed, only ensemble parallelization may be not enough for the forecast stage.

Reply to **other limitations of the proposed algorithm**:

1. *The proposed method is based on the fact that volcanic ash only occupies portion of the whole domain and hence there will be many zero rows in the ensemble state matrix. So such a method can only be implemented for flows that have such feature. It seems that few applications of Ensemble Kalman Filter (EnKF) data assimilation method have this feature.*

    Response:

    Thanks for the comment. We agree that many EnKF data assimilation applications don't have this feature. However, there are many other DA applications featuring as volcanic ash DA. We have mentioned these applications in the revision, as in **line(s) 13.29–14.3**:

    The assimilation problems where MS can achieve the acceleration effect on the computations of $\mathbf{A}^a = \mathbf{A}^f \, \mathbf{X}$ include all the volcanic-related ash/gas assimilation, e.g., assimilation of satellite data/LIDAR data/in situ data; (sand/desert) dust storm related assimilation; tornado-related assimilation; assimilation of exploding nuclear plants or factories; chemicals or oils leaking on seas; global (forecast) fire assimilation; assimilation of environmental pollutant transport, e.g., severe smog. In addition, for DA applications (e.g., Ozone, SO_2) where pollutants spread over the whole domain, usually the focus is only on the high concentrations and a threshold can be set to ignore the very low values without losing the necessary assimilation accuracy. In this case, MS can also lead to a potential acceleration since many very low concentrations can be explicitly truncated to be zeros.

2. *According to "the analysis step takes 72% of total runtime", it means the analysis step is the bottle neck. However, such case might not be general for all ash forecasts. As the computational cost for initialization and forecast greatly depends on the "forecast model" that is used. If a more expensive ash forecasting model is used, then, I would guess, the bottle neck would be ash forecasting.*

    Response:

    Agree. We have discussed on this point in the revision, as in **line(s) 16.25–16.33**:

    In this case study with the LOTOS-EUROS model (version 1.10), after the parallelization is

performed for the forecast step of EnKF assimilation, the analysis step takes 72 % of total run-time, which means the analysis step is the bottle neck. This case might not be general for all ash forecasts. As the computational cost for initialization and forecast greatly depends on the forecast model that is used. For the current development, it makes sense to use the LOTOS-EUROS model, because the model has been configured and evaluated in (Fu et al., 2015) by comparison with other famous models (e.g., NAME (Jones et al., 2007), WRF-Chem (Webley et al., 2012)) in simulating volcanic ash transport. However, if a more expensive ash forecasting model is used, then the bottle neck would be the forecast step. In this case, the forecast step should be the goal for acceleration and probably a parallel model or adaptive domain-decomposition (as discussed in Sect. 4.3) needs to be employed together with the parallel ensemble forecasts.

3. *It is not clear to me that a good parallel linear algebra library like PeTSC which allows users to specify their orderings and comes with machine optimized parallel matrix-matrix multiply operations will not outperform the versions coded up here.*

Response:

I agree that a direct parallelization (e.g., using PeTSC) for spare-dense matrix multiplication of $\mathbf{A}^f \mathbf{X}$ can probably outperform a serial MS. However, given the fact that MS can also be parallelized using similar ways or the same libraries, thus it is fair to not consider parallelization for all cases (i.e., using MS, not using MS, using CSR-based SDMM).

Actually, the parallelization in MS could be performed much easier than other approaches in dealing with $\mathbf{A}^a = \mathbf{A}^f \mathbf{X}$, because the dense-dense matrix multiplication (parallelization in the step (iv) of MS) is usually less complicated to handle than parallelization of the sparse-dense matrix multiplication (direct parallelization for $\mathbf{A}^a = \mathbf{A}^f \mathbf{X}$ or parallelized CSR-based SDMM).

In this paper, for the current usage, we keep parallelization open (a serial MS has been efficient already), since many parallelization strategies can be performed/adjusted within MS. It is for sure that the parallelization's effect on MS will be positive. However, whether a parallel version of MS is generally required (with respect to the extra efforts in parallelizing MS) depends on assimilation problems and needs to be investigated in future.

We agree that above aspects were not clear in the previous version. We have included them in **line(s) 15.28–16.8**:

Alternatively, one may also consider to (1) directly parallelize the expensive matrix multiplication of $\mathbf{A}^a_{n \times N} = \mathbf{A}^f_{n \times N} \mathbf{X}_{N \times N}$, without first performing MS; or (2) implement CSR-based SDMM (see Sect. 5) with parallelization. Both are possible alternative approaches to accelerate the expensive matrix multiplication. The first approach can be implemented by user own designed parallelization, or by utilizing scaLAPACK (`https://www.netlib.org/scalapack/`, where the main function is "pdgemm"). The second approach can be realized by using some general parallel sparse-dense matrix multiplication methods (e.g., sending each column of X and three CSR arrays of $\mathbf{A}^f$ to each processor to calculate each column of $\mathbf{A}^a$) or using a good parallel algebra library like PeTSC (`https://www.mcs.anl.gov/petsc/`) which allows users to specify own orderings and comes with machine optimized parallel matrix-matrix multiplication operations. However, given the fact that MS can also be parallelized using similar ways or same libraries, thus it is fair to not consider parallelization for all cases (i.e., using MS, not using MS, using CSR-based SDMM). Actually, the parallelization in MS could be performed much easier than other approaches in dealing with $\mathbf{A}^a = \mathbf{A}^f \mathbf{X}$, because the dense-dense matrix multiplication (parallelization in the step (iv) of MS) is easier to parallelize than the sparse-dense matrix multiplication (direct parallelization for $\mathbf{A}^a = \mathbf{A}^f \mathbf{X}$ or parallelized CSR-based SDMM).

In this paper, for the current usage, we keep the possibility of parallelization open, because a serial MS has been efficient already.

Reply to **modification suggestions and questions**:

1. *The paper focusses on "Ensemble Kalman Filter (EnKF) data assimilation method" and the new algorithm proposed in this paper is specific for "Ensemble Kalman Filter (EnKF) data assimilation method". It would be more precise to use " data assimilation based on Ensemble Kalman Filter (EnKF)" instead of just "data assimilation" in the title.*

   Response:

      Agree. As suggested, the title has been re-written in the revision,

      Accelerating volcanic ash data assimilation using a mask-state algorithm based on ensemble Kalman filter: a case study with the LOTOS-EUROS model (version 1.10)

2. *What is the subscript "q" in Eq. (3), should it be "N" ?*

   Response:

      Yes and corrected in **line(s) 4.6**.

3. *It would be interest to see the additional cost for reducing the size of the ensemble state matrix.*

   Response:

      Thanks for the comment. In the revision, we have added Table 2 to specify the computational time for all five steps of MS. And we have also added discusion related to this point, as in **line(s) 10.10–10.17**:

      The detailed experimental time for each step of MS is shown in Table 2. As expected, the dense-dense matrix multiplication in step (iv) takes the largest part (i.e., 0.8474 h for this case study) of the total computational time (0.87 h) of MS. However, step (iv) has been a big improvement compared to the case without MS (3.14 h, see Table 1), which is because the computational time for other steps (e.g., steps (i - iii) costs only 0.0156 h to reduce the size of the ensemble state matrix) is little and ignorable. Note that the total computational time of $\mathbf{A}^a = \mathbf{A}^f \mathbf{X}$ with MS (i.e., 0.87 h in Table 2) is not exactly equal to the compuational time of the MS-EnKF analysis procedures (i.e., 0.88 h in Table 1). The subtraction (i.e., 0.01 h) corresponds to the summed computational time of all the other analysis procedures (i.e., procedures 1 - 8) except for $\mathbf{A}^a = \mathbf{A}^f \mathbf{X}$ (see Fig. 2(b) and Table 3).

4. *The "non-zero rows" is around 0.393 of total rows. According to cost analysis, the cost is $O(nN^2)$. So theoretically speaking, the minimum time would be 0.393*3.14 h $\sim$ 1.234 h. Numerical test shows that the time goes down to 0.88h. It is interesting to know what are the other contributions to this time decrease. My guess is memory access cost also goes down when the size of matrix reduced.*

   Response:

      This is a good point. One reason is that when the size of matrix is reduced, the memory access cost also goes down (e.g., through better cache usages). Another possible reason is that the ash grid number actually decreases with time (not always taking 0.393 the total grid number), due to ash sedimentation and deposition processes ((Fu et al., 2015)).

      We have discussed this point in **line(s) 10.33–11.4**:

There is another interesting point. According to Fig. 3, the ash grids comprises 39.3 % of the total grids. Thus, the minimum computing time by using MS to utilize this model's characteristic should be $\approx 1.234\,\text{h}$ (i.e., $0.393 \times 3.14\,\text{h}$). However, the experimental result shows that the computational time goes down to $0.88\,\text{h}$ (see Table 1). One reason for this time decrease is that when the size of matrix is reduced, the memory access cost also goes down (e.g., through better cache usages). Another possible reason is that the ash grid number actually decreases with time (not always taking 39.3 % the total grid number), due to ash sedimentation and deposition processes ((Fu et al., 2015)).

5. *Again, regarding questions in 4,*

6. *the most expensive sub-step in the analysis step is actually a matrix multiplication. It should be trivial to parallelize it by yourself or utilize libraries like scaLAPACK. This would be a more general and more powerful way to reduce time on analysis step.*

Response to 5 and 6:

In order to speed up the matrix multiplication, we can do two things: (1) reduce the amount of operations in $\mathbf{A}^f\,\mathbf{X}$; (2) using parallel computation. For the application at hand, MS strongly reduces the total amount of computation; parallel processing can be applied after that.

In this study, although parallelization in MS is not necessarily required, we have carefully discussed its potential, in **line(s) 15.20–16.6**:

Motivated from the model's physics, the implementation MS currently is for the serial case. This implementation has reduced the computation time to an acceptable time (i.e., simulation time is less than the period of forecast in real world time). It is however interesting to discuss the potential of parallelization of the dense-dense matrix multiplication ($\tilde{\mathbf{A}}^a_{n_{ash} \times N} = \tilde{\mathbf{A}}^f_{n_{ash} \times N}\mathbf{X}_{N \times N}$) in the step (iv) of the algorithm (see Sect. 4.2 and Table 2). The related matrix multiplication can be easily parallelized on multiple processors. Optimization and evaluation on the parallelized MS will be considered in future. For the current case study, the computational time ($3.13\,\text{h}$, see Table 3) for an "ideal" reduction by parallelization of MS is not much larger than the acceleration (already) gained by MS ($2.26\,\text{h}$, subtraction between $3.13\,\text{h}$ and $0.87\,\text{h}$, see Table 3). Therefore, from application's perspective, further acceleration by parallelization is not required.

Alternatively, one may also consider to (1) directly parallelize the expensive matrix multiplication of $\mathbf{A}^a_{n \times N} = \mathbf{A}^f_{n \times N}\mathbf{X}_{N \times N}$, without first performing MS; or (2) implement CSR-based SDMM (see Sect. 5) with parallelization. Both are possible alternative approaches to accelerate the expensive matrix multiplication. The first approach can be implemented by user own designed parallelization, or by utilizing scaLAPACK (`https://www.netlib.org/scalapack/`, where the main function is "pdgemm"). The second approach can be realized by using some general parallel sparse-dense matrix multiplication methods (e.g., sending each column of X and three CSR arrays of $\mathbf{A}^f$ to each processor to calculate each column of $\mathbf{A}^a$) or using a good parallel algebra library like PeTSC (`https://www.mcs.anl.gov/petsc/`) which allows users to specify own orderings and comes with machine optimized parallel matrix-matrix multiplication operations. However, given the fact that MS can also be parallelized using similar ways or same libraries, thus it is fair to not consider parallelization for all cases (i.e., using MS, not using MS, using CSR-based SDMM). Actually, the parallelization in MS could be performed much easier than other approaches in dealing with $\mathbf{A}^a = \mathbf{A}^f\,\mathbf{X}$, because the dense-dense matrix multiplication (parallelization in the step (iv) of MS) is easier to parallelize than the sparse-dense matrix multiplication (direct parallelization for $\mathbf{A}^a = \mathbf{A}^f\,\mathbf{X}$ or parallelized CSR-based SDMM).

7. *Mathematics symbols need to be a little bit more clear, for example, in the paper, y is used as both observational space and a vector in the space?*

Response:

Agree. In the revision, we have checked all the mathematics symbols. The observation vector $\boldsymbol{y}$ is now clearly mentioned just as a vector, as in **line(s) 4.13–4.14**:

The observational network at time t_k is defined by the operator H_k which maps the state vector $\boldsymbol{x}$ to the observational vector $\boldsymbol{y}$

Accelerating volcanic ash data assimilation using a mask-state algorithm based on ensemble Kalman filter: a case study with the LOTOS-EUROS model (version 1.10)

Guangliang Fu[1], Hai Xiang Lin[1], Arnold Heemink[1], Sha Lu[1], Arjo Segers[2], Nils van Velzen[1,3], Tongchao Lu[4], and Shiming Xu[5]

[1]Delft University of Technology, Delft Institute of Applied Mathematics, Mekelweg 4, 2628 CD Delft, The Netherlands.
[2]TNO, Department of Climate, Air and Sustainability, P.O. Box 80015, 3508 TA Utrecht, The Netherlands.
[3]VORtech, P.O. Box 260, 2600 AG Delft, The Netherlands.
[4]School of Mathematics, Shandong University, Jinan, Shandong, China.
[5]Department of Earth System Science, Tsinghua University, Beijing, China

Correspondence to: Guangliang Fu (G.Fu@tudelft.nl)

Abstract. In this study, we investigate a strategy to accelerate the data assimilation (DA) algorithm. Based on evaluations of the computational time, the analysis step of the assimilation turns out to be the most expensive part. After a study on the characteristics of the ensemble ash state, we propose a mask-state algorithm which records the sparsity information of the full ensemble state matrix and transforms the full matrix into a relatively small one. This will reduce the computational cost in the analysis step. Experimental results show the mask-state algorithm significantly speeds up the analysis step. Subsequently, the total amount of computing time for volcanic ash DA is reduced to an acceptable level. The mask-state algorithm is generic and thus can be embedded in any ensemble-based DA framework. Moreover, ensemble-based DA with the mask-state algorithm is promising and flexible, because it implements exactly the standard DA without any approximation and it realizes the satisfying performance without any change of the full model.

1 Introduction

Volcanic ash erupted into atmospheres can lead to severe influences on aviation society (Gudmundsson et al., 2012). Turbine engines of airplanes are extremely threatened by the ash's ingestion (Casadevall, 1994). Thus, accurate real-time aviation advices are highly required during an explosive volcanic ash eruption (Eliasson et al., 2011). Recently, ensemble-based DA (Evensen, 2003) has been evaluated very useful to improve volcanic ash forecasts and regional aviation advices (Fu et al., 2016). It corrects volcanic ash concentrations by continuously assimilating observations. In (Fu et al., 2016), real aircraft in situ measurements were assimilated using the ensemble Kalman filter (EnKF), which is the most known and popular ensemble-based DA method. Based on the validation with independent data, ensemble-based DA was concluded powerful for improving the forecast accuracy.

However, to make the methodology efficient also in an operational (real-time) sense, the computational efforts must be acceptable. For volcanic ash DA problems, so far, no studies on the computational aspects have been reported in the literature.

1

Actually, when large amounts of volcanic ash erupted into atmospheres, the computational speed of volcanic ash forecasts is just as important as the forecast accuracy (Zehner, 2010). For example, due to the lack of a fast and accurate forecast system, the sudden eruption of the Eyjafjallajökull volcano in Iceland from 14 April to 23 May 2010, had caused an unprecedented closure of the European and North Atlantic airspace resulting in a huge global economic loss of 5 billion US dollars (Oxford-

5 Economics, 2010). Since then, research on fast and accurate volcanic ash forecasts have gained much attention, because it is needed to provide timely and accurate aviation advices for frequently operated commercial airplanes. It was shown the accuracy of volcanic ash transport can be significantly improved by the DA system in (Fu et al., 2016). Therefore, it is urgent to also consider the computational aspect, i.e., improving the computational speed of the volcanic ash DA system as fast as possible. This is the main focus of this study.

10 Due to the computational complexity of ensemble-based algorithms and the large scale of dynamical applications, applying these methods usually introduces a large computational cost. This has been reported from literature on different applications. For example, for operational weather forecasting with ensemble-based DA, Houtekamer et al. (2014) reported computational challenges at Canadian Meteorological Center with an operational EnKF featuring 192 ensemble members, using a large 600×300 global horizontal grid and 74 vertical levels. That an initialization requirement of over 7×10^{10} values to specify

15 each ensemble, results in large computational efforts on the initialization and forecast steps in weather forecasting. For oil reservoir history-matching (Tavakoli et al., 2013), the reservoir simulation model usually has a large number of state variables, thus the forecasts of an ensemble of simulation models are often time-consuming. Besides, when time-lapse seismic or dense reservoir data is available, the analysis step of assimilating these large observations becomes very time-consuming (Khairullah et al., 2013). Large computational requirements of ensemble-based DA have also been reported in ocean circulation models

20 (Keppenne, 2000; Keppenne and Rienecker, 2002), tropospheric chemistry assimilation (Miyazaki et al., 2015), and many other applications.

To accelerate an ensemble-based DA system, the ensemble forecast step can be first parallelized because the propagation of different ensemble members is independent. Thus if a computer with a sufficiently large number of parallel processors is available, all the ensemble members can be simultaneously integrated. In the analysis stage, to calculate the Kalman gain and

25 the ensemble error covariance matrix, all ensemble states must be combined together. In weather forecasting and oceanography sciences, Keppenne (2000); Keppenne and Rienecker (2002); Houtekamer and Mitchell (2001) have reported using parallelization approaches to accelerate the expensive analysis stage. In reservoir history matching, a three-level parallelization has been proposed by Tavakoli et al. (2013); Khairullah et al. (2013) in recent years, to significantly reduce computational efforts of both forecast and analysis steps due to massive dense observations and large simulation models. The first parallelization level

30 is to separately perform the ensemble simulations on different processors during the forecast step. This approach is usually quite efficient when a large ensemble size is used. However, the scale or model size of one reservoir simulation is constrained by the memory of a single processor. Thus, the second parallelization level is to perform one ensemble member simulation using a parallel reservoir model. These two levels do not deal with the analysis step, which collects all ensemble members to do computations usually on a single processor. Therefore, a third level of parallelization was implemented by Tavakoli

35 et al. (2013); Khairullah et al. (2013) through parallelizing matrix-vector multiplications in the analysis steps. Furthermore,

some other approaches on accelerating ensemble-based DA systems, have also been reported, such as GPU-based acceleration (Quinn and Abarbanel, 2011) in Numerical Weather Prediction (NWP), domain decomposition in atmospheric chemistry assimilation (Segers, 2002; Miyazaki et al., 2015). The observations used in an DA system can be also optimized with some preprocessing procedures, as reported by Houtekamer et al. (2014).

Although for other applications, there were many efforts in dealing with large computational requirements in an ensemble-based DA system, most of them cannot be directly used to accelerate volcanic ash DA. This is because the acceleration algorithms are strongly dependent on specific problems, such as model complexity (high or low resolution), observation type (dense or sparse), primary requirement (accuracy or speed). These factors determine, for a specific application, which part is the most time-consuming, and which part is intrinsically sequential. Thus, no unified approach for efficient acceleration of all the applications can be found. Although the successful approaches in other applications cannot be directly employed in volcanic ash forecasts, their success do stress the importance of designing a proper approach based on the computational analysis of a specific DA system. Therefore, the computational cost of our volcanic ash DA system will be first analyzed. Then, based on the computational analysis, we will investigate a strategy to accelerate the ensemble-based DA system for volcanic ash forecasts.

This paper is organized as follows. Sect. 2 introduces the methodology of volcanic ash DA. Sect. 3 analyzes the computational cost of the conventional volcanic ash DA system. In Sect. 4, the mask-state algorithm (MS) is developed for acceleration. The comparison between MS and standard sparse matrix methods are presented in Sect. 5. The discussions on MS is in Sect. 6. Finally, the last section summarizes the concluding remarks of our research.

2   Methodology of the volcanic ash DA system

In this study, the ensemble Kalman filter (EnKF) (Evensen, 2003) is employed to perform ensemble-based DA. EnKF is typically a sequential Monte Carlo method, according to the uncertain state estimate with N ensemble members, $\boldsymbol{\xi}_1, \boldsymbol{\xi}_2, \cdots, \boldsymbol{\xi}_N$. Each member is assumed as one sample in the distribution of the true state. It has been proposed that for operational applications, the ensemble size can be limited to 10 - 100 for cost effectiveness (Nerger and Hiller, 2013; Barbu et al., 2009). Thus, in this study, an ensemble size of 100 is used due to the high-accuracy requirement of the volcanic ash forecasts to aviation advices as mentioned in Sect. 1.

To simulate a volcanic ash plume, an atmospheric transport model is needed. In this paper, the LOTOS-EUROS (abbreviation of LOng Term Ozone Simulation - EURopean Operational Smog) model is used (Schaap et al., 2008) with model version 1.10 (http://www.lotos-euros.nl/). The LOTOS-EUROS model (Schaap et al., 2008) is an operational model focusing on nitrogen oxides, ozone, particular matter, volcanic ash. The model configurations for volcanic ash were discussed in details by Fu et al. (2016). For volcanic ash simulation, the model is configured with a state vector of size $180 \times 200 \times 18 \times 6$ (the dimensions correspond to longitude, latitude, vertical level, ash species), the size of model state is thus calculated as $\sim 10^6$.

The experiment in this study starts at t_0 (09:00 UTC, 18 May, 2010 for this study) by considering an initial condition from previous LOTOS-EUROS conventional model run (see Fig. 1(a)). In the second step (the forecast step) the model propagates

the ensemble members from the time t_{k-1} to t_k ($k > 0$, the time step is 10 minutes):

$$\boldsymbol{\xi}_j^f(k) \quad = \quad M_{k-1}(\boldsymbol{\xi}_j^a(k-1)). \tag{1}$$

The operator M_{k-1} describes the time evolution of the state which contains the ash concentrations in all model grid-boxes. The state at the time t_k has a distribution with the mean $\boldsymbol{x}^f$ and the forecast error covariance matrix $\mathbf{P}^f$ given by:

$$5 \quad \boldsymbol{x}^f(k) = [\sum_{j=1}^{N} \boldsymbol{\xi}_j^f(k)]/N, \tag{2}$$

$$\mathbf{L}^f(k) = [\boldsymbol{\xi}_1^f(k) - \boldsymbol{x}^f(k), \cdots, \boldsymbol{\xi}_N^f(k) - \boldsymbol{x}^f(k)], \tag{3}$$

$$\mathbf{P}^f(k) = [\mathbf{L}^f(k)\mathbf{L}^f(k)^T]/(N-1), \tag{4}$$

where $\mathbf{L}^f$ represents the ensemble perturbation matrix. In this study, the forecast step is performed in parallel because of the natural/common parallelism of the independent ensemble propagation, which is a trivial approach when employing ensemble-based DA (Liang et al., 2009; Tavakoli et al., 2013; Khairullah et al., 2013).

10

When the model propagates to 09:40 UTC, 18 May, 2010, the volcanic ash state gets sequentially analyzed by the DA process through combining real aircraft in situ measurements of PM_{10} and $PM_{2.5}$ concentrations until 11:10 UTC. The measurement route and values are demonstrated in Figs. 1(b, c) and the details are described in (Weber et al., 2012; Fu et al., 2016). The observational network at time t_k is defined by the operator H_k which maps the state vector $\boldsymbol{x}$ to the observational vector $\boldsymbol{y}$ by

$$15 \quad \boldsymbol{y}(k) = H_k(\boldsymbol{x}(k)) + \boldsymbol{v}(k), \tag{5}$$

where $\boldsymbol{y}$ contains the aircraft measurements and $\boldsymbol{v}$ represents the observational error. H_k selects the grid cell in $\boldsymbol{x}(k)$ that corresponds to the locations of the observation. When measurements are available, the ensemble members are updated in the analysis step using

$$\mathbf{K}(k) = \mathbf{P}^f(k)\mathbf{H}(k)^T[\mathbf{H}(k)\mathbf{P}^f(k)\mathbf{H}(k)^T + \mathbf{R}]^{-1}, \tag{6}$$

$$20 \quad \boldsymbol{\xi}_j^a(k) = \boldsymbol{\xi}_j^f(k) + \mathbf{K}(k)[\boldsymbol{y}(k) - \mathbf{H}(k)\boldsymbol{\xi}_j^f(k) + \boldsymbol{v}_j(k)], \tag{7}$$

where $\mathbf{K}$ represents the Kalman gain, $\mathbf{H}$ is the observational matrix formed by the observational operator H, $\mathbf{R}$ represents the measurement error covariance matrix and $\boldsymbol{v}_j$ represents the realization out of the observation error distribution $\boldsymbol{v}$. After the continuous assimilation ending at 11:10 UTC, the forecast at 12:00 UTC is illustrated in Fig. 1(d), for which the forecast accuracy has been carefully evaluated as significantly improved compared to the case without DA (Fu et al., 2016).

25

The EnKF with above setups is abbreviated as "conventional EnKF" and used in this study for the computational evaluation. Note that in the study we don't use covariance localization as proposed by Hamill et al. (2001) for reducing spurious covariances. This is because although localization is possible, the ideal case is not to use it in order to have the correct covariances in a large (converged) ensemble. It is crucial for localization that when unphysical (spurious) covariances are eliminated, physical (correct) covariances can be well maintained (Petrie and Dance, 2010). If the "filtering length scale" for localization is too

30

long (i.e., all the dynamical covariances are allowed), many of the spurious covariances may not be eliminated. If the length is

too short, important physical dynamical covariances then may be lost together with the spurious ones. Therefore, essentially deciding an accurate localization is a challenging subject (Riishojgaard, 1998; Kalnay et al., 2012) especially for accuracy-demanding applications. Therefore, in this study we choose the ensemble size of 100 to guarantee the accuracy and avoid large spurious covariances.

5 ## 3 Computational analysis for volcanic ash DA

3.1 Computational analysis of the total runtime

Ensemble-based DA is a useful approach to improve the forecast accuracy of volcanic ash transport. However, if it is time-consuming, it cannot be taken as efficient due to the high requirement on speed for volcanic ash DA (see Sect. 1). Based on this consideration, we need to analyze the computational cost of a conventional volcanic ash DA system.

10   As introduced in Sect. 2, the total execution time of conventional EnKF comprises four parts, i.e., initialization, forecast, analysis and other computational cost. The initialization time includes reading meteorological data, initializing model geographical and grid configurations, reading emission information, initializing stochastic observer for reading and transforming observations to the model grid, initializing all the ensemble states and ensemble mean, and so on. The forecast time is obtained from Eq. (1), while the analysis time corresponds to the computational sum from Eq. (2) to (7). The other computational time

15 includes script compiling, setting environment variables, starting and finalizing DA algorithms, etc.

The evaluation result of the conventional EnKF is shown in Table 1 (the middle column). It can be seen that the total computational time (4.36 h) is relatively large compared to the simulation window (3.0 h, i.e., from 9:00 - 12:00 UTC, 18 May, 2010), which is too much in an operational sense. Therefore, in this study, we aim to accelerate the computation to within an acceptable runtime (i.e., requires less runtime than the time period of the DA application).

20   It can be also observed from Table 1 that the main contribution to the total execution time is the analysis step. Compared to the initialization and forecast time, the analysis stage takes 72 % of the total runtime. Due to the expensive analysis step, although some approaches (such as MPI-parallel I/O (Filgueira et al., 2014), domain decomposition (Segers, 2002)) can potentially accelerate the initialization and forecast step, the effect to the final acceleration of the total computational cost is little. Therefore, to get acceptable computational time, the cost reduction in the analysis step is the target. One may wonder that since

25 the amount of observations is small, why does analysis takes so much time? The large state vector seems to be left responsible for the problem. To know the exact reason, the detailed computational cost of the analysis step must be evaluated.

3.2 Cost estimation of all analysis procedures

We start with the formulations of the analysis step. The analysis step is represented by Eq. (7), which can be written in a full matrix format with Eq. (8),

30 $$\mathbf{A}_{n \times N}^{a} = \mathbf{A}_{n \times N}^{f} + \mathbf{K}_{n \times m}(\mathbf{Y}_{m \times N} - \mathbf{H}_{m \times n}\mathbf{A}_{n \times N}^{f}), \tag{8}$$

where the subscripts represent the matrix's dimensions. $\mathbf{A}^f$ and $\mathbf{A}^a$ represent the forecasted and analyzed ensemble state matrix, and are respectively built up from $\boldsymbol{\xi}^f$ and $\boldsymbol{\xi}^a$ with N ensembles. The measurement ensemble matrix $\mathbf{Y}$ is formed by an ensemble of $\boldsymbol{y} + \boldsymbol{v}$ (see Eq. (7)). $\mathbf{H}$ is the observational matrix, which is used to select state variables (at measurement locations) in the full ensemble state matrix corresponding to the measurement ensemble matrix $\mathbf{Y}$. n is the number of model
5 state variables in a three-dimensional (3-D) domain, i.e., $\sim 10^6$ in this study (see Sect. 2). m is the amount of measurements at one assimilation time, which depends on the measurement type. For aircraft in situ measurements used in this study (see Fig. 1(c)), two measurements are made at each time by one research flight, so that m is 2 here. N is the ensemble size and is taken as 100 in this study. As described in Eq. (3), the ensemble perturbation matrix $\mathbf{L}^f$ in EnKF can be re-written as

$$\mathbf{L}^f_{n \times N} = \mathbf{A}^f_{n \times N} - \bar{\mathbf{A}}^f_{n \times N} = \mathbf{A}^f_{n \times N}(\mathbf{I}_{N \times N} - \frac{1}{N}\mathbf{1}_{N \times N}) = \mathbf{A}^f_{n \times N}\mathbf{B}_{N \times N}, \tag{9}$$

10 where $\mathbf{I}$ is an $N \times N$ unit matrix and $\mathbf{1}$ is an $N \times N$ matrix with all elements equal to 1. Thus, $\mathbf{L}^f = \mathbf{A}^f\mathbf{B}$ where $\mathbf{B}_{N \times N}$ is introduced to represent $(\mathbf{I}_{N \times N} - \frac{1}{N}\mathbf{1}_{N \times N})$. So that, $\mathbf{HL}^f = \mathbf{O}^f\mathbf{B}$, where $\mathbf{O}^f_{m \times N}$ is used to represent $(\mathbf{HA}^f)$. Here we explicitly express $\mathbf{L}^f$ and $\mathbf{HL}^f$ in the form of $\mathbf{A}^f$ and $\mathbf{O}^f$, respectively. This is because in our volcanic ash DA system, $\mathbf{A}^f$ and $\mathbf{O}^f$ are two of the three inputs (another one is the measurement ensemble matrix $\mathbf{Y}$ for the analysis step. These are the three inputs used for actual computations in the analysis step. As shown in Fig. 2(a), $\mathbf{A}^f$ is obtained from the forecast step,
15 $\mathbf{O}^f$ and $\mathbf{Y}$ are acquired from our stochastic observer module (see Fig. 2(a)) which is used for a volcanic ash transport model to integrate geophysical measurements. With the input $\mathbf{Y}$, the measurement error covariance $\mathbf{R}$, as introduced in Eq. (6), can be then computed with

$$\mathbf{R}_{m \times m} = \frac{1}{N-1}(\mathbf{Y}_{m \times N} - \bar{\mathbf{Y}}_{m \times N})(\mathbf{Y}_{m \times N} - \bar{\mathbf{Y}}_{m \times N})' = \frac{1}{N-1}(\mathbf{YB})(\mathbf{YB})'. \tag{10}$$

Based on previous definitions and Eqs. (2) to (7), the analysis step can be reformulated as followings,

20 $$
\begin{aligned}
\mathbf{A}^a_{n \times N} &= \mathbf{A}^f + \mathbf{K}(\mathbf{Y} - \mathbf{HA}^f) \\
&= \mathbf{A}^f + \mathbf{P}^f\mathbf{H}'(\mathbf{HP}^f\mathbf{H}' + \mathbf{R})^{-1}(\mathbf{Y} - \mathbf{HA}^f) \\
&= \mathbf{A}^f + \frac{1}{N-1}\mathbf{L}^f(\mathbf{HL}^f)'[\frac{1}{N-1}(\mathbf{HL}^f)(\mathbf{HL}^f)' + \frac{1}{N-1}(\mathbf{YB})(\mathbf{YB})']^{-1}(\mathbf{Y} - \mathbf{HA}^f) \\
&= \mathbf{A}^f + \mathbf{A}^f\mathbf{B}(\mathbf{O}^f\mathbf{B})'[(\mathbf{O}^f\mathbf{B})(\mathbf{O}^f\mathbf{B})' + (\mathbf{YB})(\mathbf{YB})']^{-1}(\mathbf{Y} - \mathbf{O}^f) \\
&= \mathbf{A}^f\{\mathbf{I} + \mathbf{B}(\mathbf{O}^f\mathbf{B})'[(\mathbf{O}^f\mathbf{B})(\mathbf{O}^f\mathbf{B})' + (\mathbf{YB})(\mathbf{YB})']^{-1}(\mathbf{Y} - \mathbf{O}^f)\} \\
\end{aligned}
\tag{11}
$$
25 $$= \mathbf{A}^f_{n \times N}\mathbf{X}_{N \times N} \quad,$$

where

$$\mathbf{X}_{N \times N} = \{\mathbf{I} + \mathbf{B}(\mathbf{O}^f\mathbf{B})'[(\mathbf{O}^f\mathbf{B})(\mathbf{O}^f\mathbf{B})' + (\mathbf{YB})(\mathbf{YB})']^{-1}(\mathbf{Y} - \mathbf{O}^f)\}. \tag{12}$$

Eq. (11) shows how the analysis step is performed in a volcanic ash DA system. In order to accelerate the analysis step, the most time-consuming part must be reduced. Fig. 2(b) shows estimations of the computational cost for each procedure in the analysis
30 step. Considering that the state number n ($\sim 10^6$) is significantly larger than the measurement number m (m=2 here) and the

ensemble size N (N=100), thus the most time-consuming procedure in the analysis step is the last one, that is $\mathbf{A}^a = \mathbf{A}^f \mathbf{X}$ with computational cost of O($n N^2$). Therefore, in our volcanic ash DA system, this part is the most time-consuming part in the analysis step. Note that the procedure $[(\mathbf{O}^f \mathbf{B})(\mathbf{O}^f \mathbf{B})' + (\mathbf{YB})(\mathbf{YB})']^{-1}$ for Singular Value Decomposition (SVD) in our study is not time-consuming, which is different from applications of reservoir history matching (Tavakoli et al., 2013; 5 Khairullah et al., 2013). This is because the SVD procedure costs O(m^3), and due to the measurement size in the order of the size of the state in those cases, SVD procedure thus requires a huge computational cost for reservoir DA.

4 The mask-state algorithm (MS) for acceleration of the analysis step

4.1 Characteristic of ensemble state matrix $\mathbf{A}^f$

Analysis in the previous section shows that $\mathbf{A}^a = \mathbf{A}^f \mathbf{X}$ is most expensive in the analysis step. Each column of $\mathbf{A}^f$ is con-
10 structed from a forecasted ensemble state, thus the dimension of $\mathbf{A}^f$ is $n \times N$. In each column, the element values correspond to volcanic ash concentrations in a 3-D domain. Fig. 3 shows the coverage of all ensemble forecast states at a selected time 10:00 UTC 18 May, 2010, without loss of generality. A common phenomenon can be observed, that is only a part of the 3-D domain are filled with volcanic ash. The ash clouds only concentrate in a plume which is transported over time. This is because volcanic eruption is a fast and strong process. The advection dominates the transport, and the volcanic ash plume is transported
15 with the wind. This is a particular characteristic for volcanic ash transport, in contrast to other atmospheric related applications such as ozone (Curier et al., 2012), SO_2 (Barbu et al., 2009), CO_2 (Chatterjee et al., 2012). For those applications, the concentrations are everywhere in the domain, the emission sources are also everywhere, and observations are available throughout the domain too (especially for satellite data). Whereas for application of volcanic ash transport, the source emission is only at the volcano, thus usually only a limited domain is polluted by ash. As shown in Fig. 3, in the 3-D domain with grid size of
20 3.888×10^6, the number of grids in the area with volcanic ash is counted as 1.528×10^6, whereas the number of no-ash grids is 2.36×10^6. Note that shown in the figure are accumulated ash coverages of all ensemble states, thus in the no-ash grids, there are no ash for all the ensemble states. Thus a very large number of rows in $\mathbf{A}^f$ are zero corresponding to the no-ash grids. These zero rows in $\mathbf{A}^f$ have no contributions to $\mathbf{A}^a = \mathbf{A}^f \mathbf{X}$, because a zero row in $\mathbf{A}^f$ always results in a zero row in $\mathbf{A}^a$. Therefore, for the case of Fig. 3, $\frac{2}{3}$ of the computations are redundant and can be avoided. To realize this, one may
25 think to limit the domain for the entire assimilation steps, then the number of zero rows certainly would be largely reduced. This is actually incorrect, because these zero rows are changing along with the transport of ash clouds, and not constant at each analysis step. So the full domain must be considered and it should be adaptive (choose different zero rows according to different $\mathbf{A}^f$ at different analysis time).

4.2 Derivation of the mask-state algorithm (MS)

30 Here we introduce item n_{noash} to represent the number of zero rows in the ensemble state matrix $\mathbf{A}^f$, and use n_{ash} to represent the number of other rows (also n_{ash} represents the grid size of ash plume). When computing $\mathbf{A}^a = \mathbf{A}^f \mathbf{X}$, to avoid

all the computations related to n_{noash} rows with zero elements, the index of other n_{ash} rows must be first decided. This index is meant to reduce the dimensions of $\mathbf{A}^f$. After getting a $\mathbf{A}^a$ with a dimension of $n_{ash} \times N$, the index will be used again to reconstruct the full matrix $\mathbf{A}^a$ with the dimension of $n \times N$. Based on this idea, we propose a mask-state algorithm (MS) which deals with the time-consuming analysis update. MS includes five steps:

5 (i) **Compute ensemble mean state $\bar{\mathbf{A}}^f$**: The mean state $\bar{\mathbf{A}}^f_{n \times 1}$ can be easily computed by averaging $\mathbf{A}^f_{n \times N}$ along N columns. Due to all elements in $\mathbf{A}^f_{n \times N}$ corresponding to ash concentrations, thus all elements in $\mathbf{A}^f_{n \times N}$ are larger than or equal to zero. The index of non-zero rows in $\bar{\mathbf{A}}^f_{n \times 1}$ is thus equivalent to that in $\mathbf{A}^f_{n \times N}$. The computational cost for this step is $\mathrm{O}(nN)$.

(ii) **Construct mask array $\mathbf{z}$**: Based on previously obtained $\bar{\mathbf{A}}^f_{n \times 1}$, we search the non-zero elements of $\bar{\mathbf{A}}^f_{n \times 1}$ and record

10 the index into a mask array $\mathbf{z}_{n_{ash} \times 1}$. With this strategy, we don't need to search the full matrix $\mathbf{A}^f_{n \times N}$ and build an index matrix for storage. This is a benefit for saving memory. The computational cost for this step is $\mathrm{O}(n)$.

(iii) **Construct masked ensemble state matrix $\tilde{\mathbf{A}}^f$**: Using the mask array $\mathbf{z}_{n_{ash} \times 1}$ obtained from step (ii), $\tilde{\mathbf{A}}^f_{n_{ash} \times N}$ can be constructed column by column according to Eq. (13), and the computational cost (overhead) for this step is $\mathrm{O}(n_{ash} N)$.

$$\tilde{\mathbf{A}}^f(1:n_{ash}, 1:N) = \mathbf{A}^f(\mathbf{z}(1:n_{ash}), 1:N), \tag{13}$$

15 (iv) **Compute $\tilde{\mathbf{A}}^a$ by multiplying $\tilde{\mathbf{A}}^f$ and $\mathbf{X}$**: Perform matrix computation $\tilde{\mathbf{A}}^a_{n_{ash} \times N} = \tilde{\mathbf{A}}^f_{n_{ash} \times N} \mathbf{X}_{N \times N}$. This step is similar to $\mathbf{A}^a = \mathbf{A}^f \mathbf{X}$, as described in Sect. 3.2, but the computational cost now becomes $\mathrm{O}(n_{ash} N^2)$ instead of $\mathrm{O}(n N^2)$.

(v) **Construct analyzed ensemble state matrix $\mathbf{A}^a$**: With the computed $\tilde{\mathbf{A}}^a$ from step (iv) and the mask array $\mathbf{z}$ from step (ii), the final analyzed ensemble state matrix $\mathbf{A}^a_{n \times N}$ can be constructed based on Eq. (14). The computational cost (overhead) for this step is $\mathrm{O}(nN)$.

20 $$\mathbf{A}^a(\mathbf{z}(1:n_{ash}), 1:N) = \tilde{\mathbf{A}}^a(1:n_{ash}, 1:N), \tag{14}$$

According to the derivations of MS, the computational cost related to zero rows are avoided. Here the "zero rows" doesn't equal to "zero elements". The former corresponds to the regions where there are no ash for all the ensemble members, while the latter also counts the no-ash regions specifically for some ensembles. Certainly the consideration of all "zero elements" can include all the sparsity information of the ensemble state matrix, but extra computations and memories must be spent on

25 searching the full matrix $\mathbf{A}^f_{n \times N}$ with a computational cost of $\mathrm{O}(nN)$ and storing a mask state matrix with dimensions of $n \times N$. This is expensive compared to construct the mask array in the procedure (ii). Actually, after a careful check on the volcanic ash ensemble plumes, there is no "bad" ensemble which is really different from others. Although the concentration level in ensemble members are distinct, the main direction and the occurrence to the grid cells are more or less same. This means, the "zero rows" actually more or less equals to "zero elements", but much faster than the way with "zero elements",

30 which confirms the suitability and advantage of procedure (ii). Probably when there are big meteorological uncertainties, the

"zero elements" will be much larger than "zero rows". In this case, how to make use of the sparsity information in the ensemble state matrix, will be considered in future.

Based on procedures of MS, the computational cost of $\mathbf{A}^a = \mathbf{A}^f \mathbf{X}$ can be reduced. However, without a careful evaluation, we cannot conclude MS is fast, because the algorithm also employs other procedures. If these procedures (i)(ii) (iii)(v) are

5   much cheaper than the main procedure (iv), MS can definitely speed up the analysis step, and vice versa. Now we analyze MS's computational cost, which can be summed as $\mathrm{O}(n\,N)$+ $\mathrm{O}(n)$+ $\mathrm{O}(n_{ash}\,N)$+ $\mathrm{O}(n_{ash}\,N^2)$+ $\mathrm{O}(n\,N)$, i.e., $\mathrm{O}(n\,N + n_{ash}\,N^2)$. Thus, the computational overhead involved to transform the full matrix to a small one (i.e., $\mathrm{O}(n_{ash}\,N)$ for procedure (iii)) has little effect in the total computation cost of MS (i.e., $\mathrm{O}(n\,N + n_{ash}\,N^2)$). However, the computational overhead of transforming the small matrix to the full one (i.e., $\mathrm{O}(n\,N)$ for procedure (v)) does contribute a part, which cannot be ignored, to the total

10   MS's computational cost. The computational cost without MS is $\mathrm{O}(n\,N^2)$.

The comparison between both cost (with and without MS, i.e., $\mathrm{O}(n\,N + n_{ash}\,N^2)$ and $\mathrm{O}(n\,N)$) indicates when the number of non-zero rows (n_{ash}, i.e., the number of grids with ash) of the forecasted ensemble state matrix satisfies $n_{ash} < \frac{N-1}{N}n$, then MS can accelerate $\mathbf{A}^a = \mathbf{A}^f \mathbf{X}$. Here, $\mathrm{O}(n\,N + n_{ash}\,N^2)$ and $\mathrm{O}(n\,N)$ are of the same order when $n_{ash} < \frac{N-1}{N}\,n$. The larger the difference between n_{ash} and $\frac{N-1}{N}n$, the better the speedup can be achieved. According to this analysis, and the

15   characteristic (e.g., $\frac{n_{ash}}{n}$ approximately equals to $\frac{1}{3}$ in this case) of volcanic ash transport as described in Sect. 4.1, the relation is certainly satisfied and is actually $n_{ash} \ll \frac{N-1}{N}\,n$ (significantly smaller) for our study. Therefore, for our volcanic ash DA system, with MS, the computational cost for the time-consuming part $\mathbf{A}^a = \mathbf{A}^f \mathbf{X}$ is $\mathrm{O}(n_{ash}\,N^2)$, which is much reduced compared to $\mathrm{O}(n\,N^2)$ with conventional computations.

The relation $n_{ash} < \frac{N-1}{N}\,n$ indicates whether we would have speedup by the MS method, actually it can be extended to Eq.

20   (15),

$$S_{\mathrm{ms}} = \frac{\mathrm{O}(n\,N^2)}{\mathrm{O}(n\,N + n_{ash}\,N^2)} = \mathrm{O}(\frac{n}{n_{ash}}), \tag{15}$$

which explicitly specifies the expected amount of speedup (S_{ms}) of $\mathbf{A}^a = \mathbf{A}^f \mathbf{X}$ by the MS algorithm. In this case study, N is taken at 100 and $\frac{n_{ash}}{n} \approx \frac{1}{3}$, so S_{ms} is approximately 3.0.

According to Amdahl's law (Amdahl, 1967), the total computational speedup (S_{total}) by MS can be predicted by Eq. (16),

25
$$S_{\mathrm{total}} = \frac{1}{(1 - p_{\mathrm{ms}}) + \frac{p_{\mathrm{ms}}}{S_{\mathrm{ms}}}}, \tag{16}$$

where p_{ms} is the proportion of the computational cost of $\mathbf{A}^a = \mathbf{A}^f \mathbf{X}$ in the overall DA computations. It has been evaluated that the computational cost of $\mathbf{A}^a = \mathbf{A}^f \mathbf{X}$ dominates the analysis step (see Fig. 2(b)), thus the proportion of the computational cost of $\mathbf{A}^a = \mathbf{A}^f \mathbf{X}$ approximates the proportion of the analysis step in the total DA computations (i.e., $p_{\mathrm{ms}} \approx 72\,\%$ in this case, as described in Sect. 3.1). Therefore, based on Eq. (16), the maximum ("ideal") computational speedup can be predicted

30   to be $\frac{1}{1 - p_{\mathrm{ms}}}$ (i.e., ≈ 3.57 for this case study) when S_{ms} approximates infinity. However, this is not the actual speedup because S_{ms} is in fact specified by Eq. (15). Based on discussions above, S_{total} can be therefore estimated by Eq. (15) at ≈ 2.0 in this case.

4.3 Experimental results

Analysis of the algorithmic complexity of MS shows that MS is an efficient approach to reduce the computational cost of the time-consuming $\mathbf{A}^a = \mathbf{A}^f \mathbf{X}$. Now MS will be applied in the real volcanic ash DA system, to investigate whether in practice it can well speed up the analysis step. We perform MS in the conventional EnKF, which means initialization, forecast steps are all computed as the conventional EnKF. The only difference between MS-EnKF and conventional EnKF is that in the former MS is employed for analysis step, and in the latter is standard analysis step. The result and related specifications are shown in Table 1. As introduced in Sect. 2, the forecast step has been configured with the conventional parallelization, thus N+2 (102 here) cores are actually used (one core for the DA algorithm, the other N+1 cores for the parallel forecast of N ensemble members and one ensemble mean). It can be seen from Table 1 that MS indeed largely accelerates the analysis step (as expected, by a factor of about 3.0 for this study) which confirms the theoretical cost evaluation. The detailed experimental time for each step of MS is shown in Table 2. As expected, the dense-dense matrix multiplication in step (iv) takes the largest part (i.e., 0.8474 h for this case study) of the total computational time (0.87 h) of MS. However, step (iv) has been a big improvement compared to the case without MS (3.14 h, see Table 1), which is because the computational time for other steps (e.g., steps (i - iii) costs only 0.0156 h to reduce the size of the ensemble state matrix) is little and ignorable. Note that the total computational time of $\mathbf{A}^a = \mathbf{A}^f \mathbf{X}$ with MS (i.e., 0.87 h in Table 2) is not exactly equal to the compuational time of the MS-EnKF analysis procedures (i.e., 0.88 h in Table 1). The subtraction (i.e., 0.01 h) corresponds to the summed computational time of all the other analysis procedures (i.e., procedures 1 - 8) except for $\mathbf{A}^a = \mathbf{A}^f \mathbf{X}$ (see Fig. 2(b) and Table 3).

MS is now experimentally proven as efficient to significantly reduce the computational time for the analysis step during volcanic ash DA. Note that it can also be observed that the computational time for the "other" parts in Table 1 (such as operations for setting environmental variables, starting and finalizing DA algorithms, as mentioned in Sect. 3.1) is slightly reduced by the MS method (i.e., 0.03 h in this case). This is because in the conventional EnKF, the ensemble mean state $\bar{\mathbf{A}}^f$ is calculated in the "other" parts as an output to finalize the DA algorithms, while in MS-EnKF, the calculations of $\bar{\mathbf{A}}^f$ are needed and directly involved in the "Analysis" part.

The result shows that benefiting from the success of reduced analysis step, the overall computational cost indeed gets significantly reduced. The total execution time is 1.95 h which is less than the simulation window of 3 h (09:00 - 12:00 UTC, May 18, 2010). This result satisfies our goal to accelerate the computation to an acceptable runtime (i.e., requires less run time than the time period of the DA application). Therefore, aviation advices based on the MS-EnKF can be provided as not only accurate, but also sufficiently fast. Note that the result (1.95 h) is obtained after the volcanic ash is transported to the continental Europe. If the assimilation is performed in the starting phase of volcanic ash eruption (when aircraft measurements are available), a more significant acceleration would be obtained. This is because in this case the volcanic ash is only transported in an area near to the volcano, thus the number of no-ash grid cells will take a large proportion (much higher than $\frac{2}{3}$ for this case study) of the full domain.

There is another interesting point. According to Fig. 3, the ash grids comprises 39.3 % of the total grids. Thus, the minimum computing time by using MS to utilize this model's characteristic should be ≈ 1.234 h (i.e., 0.393×3.14 h). However, the

experimental result shows that the computational time goes down to 0.88 h (see Table 1). One reason for this time decrease is that when the size of matrix is reduced, the memory access cost also goes down (e.g., through better cache usages). Another possible reason is that the ash grid number actually decreases with time (not always taking 39.3 % the total grid number), due to ash sedimentation and deposition processes ((Fu et al., 2015)).

5   Note that in this study, we only perform the commonly used ensemble parallelization for the forecast step (already efficient compared to the expensive analysis step), but do not choose model-based parallelization (e.g., tracer or domain decomposition). As specified in Table 1, no parallelization is implemented on the 6 tracers. This is because due to the important aggregation process (Folch et al., 2010), there are big dependencies between different ash components and thus it doesn't make much sense to parallelize them. As for domain-decomposed parallelization (Segers, 2002), it is not efficient for our application.

10 This is because volcanic ash is special in the sense that the model is only doing computations in a small part of the domain (i.e., there is no data in a rather large part of domain), and this active part is continuously changing. Thus, a fixed domain decomposition is not very useful here because of the changing plume position. In this sense, some advanced approach such as adaptive domain-decomposed parallelization (Lin et al., 1998) should be adopted to achieve additional acceleration to the volcanic ash forecast stage. This is an interesting subject for future application, when a more complicated model is employed,

15 only ensemble parallelization may be not enough for the forecast stage.

5   Comparison between MS and standard sparse matrix methods

5.1   Issues related to the generation of CSR-based arrays

According to Sect. 4, MS has proven to be capable of solving the computational issue of $\mathbf{A}^a = \mathbf{A}^f \mathbf{X}$. Motivated from model's characteristics, MS was proposed from an application's perspective and achieved a good result by managing the irregular

20 sparsity in our complicated volcanic ash DA system. The main reason why MS is efficient is that the sparsity of $\mathbf{A}^f$ can be well utilized by MS. In Sect. 4, we only performed the comparison between MS and the case of full storage dense matrices. However, the problem abstracted here ($\mathbf{A}^f \mathbf{X}$) is actually a sparse-dense matrix multiplication (SDMM) problem, since $\mathbf{A}^f$ is sparse and $\mathbf{X}$ is dense (see Eq. (12) for $\mathbf{X}$). Thus, one may wonder what the result would be, if the comparison of MS is made to more standard sparse matrix methods, such as Compressed Sparse Row (CSR)-based methods (Saad, 2003; Bank and

25 Douglas, 1993), which are commonly used for sparse matrix vector/matrix multiplication.

Before we make the comparison, we need to first address the intrinsic problem when considering standard sparse matrix methods in EnKF for $\mathbf{A}^a = \mathbf{A}^f \mathbf{X}$. The issue is that it is not possible to directly generate a sparse storage format (e.g., CSR) of $\mathbf{A}^f$ without first generating the full matrix $\mathbf{A}^f$. This is mainly because $\mathbf{A}^f$ comes from the model-driven ensemble forecast step, where each $\mathbf{A}^f$ column corresponds to one member of the ensemble. During model forecast, we know there are indeed

30 no-ash grids. However, it is not sure where the plume exactly is after one forecasting time step. This is highly dependent on the weather conditions and the model processes (e.g, advection and diffusion for horizontal grids, sedimentation and deposition for vertical grids). Thus, a fixed and wide domain is usually needed by the model to avoid complications, resulting in the generation of the full storage of $\mathbf{A}^f$ (to be used in $\mathbf{A}^a = \mathbf{A}^f \mathbf{X}$). Therefore, if we want to implement a CSR storage format for

the sparse matrix $\mathbf{A}^f$, we must first generate the full storage $\mathbf{A}^f$ from the ensemble forecast step, and then we generate the three CSR arrays based on $\mathbf{A}^f$.

Generating CSR arrays is usually much more expensive (computationally) than a single sparse matrix-vector multiplication (SpMV). Thus, if we generate CSR arrays for only performing one-time SpMV, it would be meaningless from HPC's point of view. Fortunately, this is not the case for $\mathbf{A}^a = \mathbf{A}^f \mathbf{X}$ (i.e., SDMM), which can actually be considered as N-times SpMV. (Here, $\mathbf{X}$ has N columns, and one SpMV means the multiplication of $\mathbf{A}^f$ with one column of $\mathbf{X}$.) Thus, CSR-based SDMM might also be a candidate in reducing the computation time of $\mathbf{A}^a = \mathbf{A}^f \mathbf{X}$. It remains interesting to compare the performance of CSR-based SDMM and MS, in dealing with $\mathbf{A}^a = \mathbf{A}^f \mathbf{X}$ for our study case.

5.2 Result of CSR-based SDMM

To implement CSR-based SDMM for $\mathbf{A}^a = \mathbf{A}^f \mathbf{X}$, the three CSR arrays for $\mathbf{A}^f$ (denoted $\boldsymbol{val}$, $\boldsymbol{col\_idx}$, $\boldsymbol{row\_ptr}$ in this study) need to be first generated. The array $\boldsymbol{val}$ of size $n_{\boldsymbol{val}}$ stores non-zero values of $\mathbf{A}^f$, where $n_{\boldsymbol{val}} \approx n_{ash} N$. (In this study case, $n = 3.888 \times 10^6$, $n_{ash} \approx \frac{1}{3} n$, $N = 100$.) The array $\boldsymbol{col\_idx}$ of the same size $n_{\boldsymbol{val}}$ stores column index of the non-zeros. The array $\boldsymbol{row\_ptr}$ saves the start and the end pointers of the non-zeros of the rows in $\mathbf{A}^f$. The size of $\boldsymbol{row\_ptr}$ is $n + 1$.

After the above three CSR arrays are generated, CSR-based SpMV can be performed for multiplying $\mathbf{A}^f$ with a column vector $\boldsymbol{v}$ in $\mathbf{X}$ (see Algorithm 1 in Fig. 4(a)). After that, Algorithm 2 (Fig. 4(b)) can be implemented by looping Algorithm 1 for N times to obtain $\mathbf{A}^a = \mathbf{A}^f \mathbf{X}$. The experimental result of CSR-based SDMM is shown in Table 4, where all the environmental conditions (such as the DA system, the programming environment) are the same as the case of MS. This gives a fair comparison between CSR-based SDMM and MS. In addition, for a pure algorithmic comparison with the serial MS, here the CSR-based SDMM is also performed in a serial case.

From Table 4, we can first confirm that the computational time (i.e., 0.0407 h) for the generation of the three CSR-based arrays ($\boldsymbol{val}$, $\boldsymbol{col\_idx}$, $\boldsymbol{row\_ptr}$ to represent the sparse matrix $\mathbf{A}^f$) indeed takes more time than the computational time of one CSR-based SpMV (i.e., 0.0117 h). Thus, there is little value to perform sub-step (i) (see Table 4) if only one SpMV (i.e., sub-step (ii)) is needed. However, to get all N (i.e., 100) columns of $\mathbf{A}^a$, the sub-step (ii) is looped for N times, resulting in an ignorable impact of sub-step (i) on the total computational time (i.e., 1.21 h) of CSR-based SDMM.

The result of CSR-based SDMM also shows that the standard sparse matrix methods can reduce the computational time of $\mathbf{A}^a = \mathbf{A}^f \mathbf{X}$, by comparing with the conventional way in Table 3. However, it can also be observed that the computational time of CSR-based SDMM is larger than MS (i.e., 1.21 h versus 0.87 h in Table 3). Thus, although application of sparse matrix multiplication methods is positive, it is still slower than MS on our problem.

5.3 Comparison between CSR-based SDMM and MS

In the CSR-based SDMM, only non-zero elements in $\mathbf{A}^f$ participate in the multiplication between $\mathbf{A}^f$ and $\mathbf{X}$, thus redundant computation (related to zero elements in $\mathbf{A}^f$) is avoided. So the computation time of $\mathbf{A}^f \mathbf{X}$ is reduced with CSR-based SDMM. In the following, we analyze the performance difference between CSR-based SDMM and MS.

Firstly, from the programming's perspective, in CSR-based SDMM, the loop number for the rows of $\mathbf{A}^f$ is from 1 to n (see Fig. 4(a)), while the corresponding loop number in MS is from 1 to n_{ash} (see the step (iv) of MS in Sect. 4.2, $n_{ash} \approx \frac{1}{3}n$). Although only non-zero elements are used in the multiplication in CSR-based SDMM, the length of the outer loop is still n (much larger than n_{ash}) which is the essential reason that MS is faster than CSR-based SDMM. Note that as discussed in Sect. 4.1, there are many zero rows in $\mathbf{A}^f$, thus CSR-based SDMM actually does nothing when it comes to a zero row, but still needs to execute the loop. Within each loop number, it has to check the information from $\boldsymbol{row\_ptr}$ (size $n+1$), where the value corresponding to a zero row is usually set to be the value in $\boldsymbol{row\_ptr}$ corresponding to the first following non-zero row.

Secondly, with respect to the algorithm, CSR-based SDMM utilizes the sparsity of $\mathbf{A}^f$ by its generation of three CSR arrays. While, MS not only utilizes the sparsity information of the sparse matrix $\mathbf{A}^f$, but also utilizes the consistency of ensemble forecasts, that is, ensemble forecasted states are not consistent in values but usually consistent in non-zero locations. This is a typical property in ensemble-based DA, resulting in that N ensemble plumes are different in concentration values, but have similar transport directions/shapes (see Sect. 4.2). Thus, most of the zero elements in $\mathbf{A}^f$ are actually in zero rows of $\mathbf{A}^f$ for an EnKF application, which leads to a small number of non-zero rows (n_{ash}) compared to the full number of rows (n) of $\mathbf{A}^f$. Therefore, the way only considering n_{ash} rows in $\tilde{\mathbf{A}}^f_{n_{ash} \times N} \mathbf{X}_{N \times N}$ (see the step (iv) of MS in Sect. 4.2) is more advantageous for an EnKF application than that considering all n rows in CSR-based SDMM. Based on the above analysis, MS can be considered as a specific sparse matrix method, which typically works for ensemble-based DA applications.

It is useful to apply standard sparse matrix methods (e.g., CSR-based SDMM) for our assimilation application. The accelerated analysis step by CSR-based SDMM (1.22 h, see Table 3) also reduces the total computational time (i.e., 2.29 h, see Table 1 for the computational time of initialization, forecast and others) to an acceptable level (i.e., less than 3 h for our case study). In practice, due to the better performance of MS than CSR-based SDMM, we will use MS as a better choice for assimilation applications. In addition, we do not only intend to present MS, but also intend to reveal which part is the most time-consuming part for plume-type assimilation of in situ observations.

6 Discussions on MS

6.1 Applicability

For volcanic ash forecasts, only a relatively small domain is polluted compared to the full 3-D domain, so that MS can work efficiently. Using MS is also applicable for many other DA problems, where the domain is not fully polluted by the species. It does not matter what the emission looks like and whether the releases are short- or long-lived species. Given an assimilation problem, the only restriction for MS to gain an acceleration is whether the whole domain is fully polluted or partly polluted. The assimilation problems where MS can achieve the acceleration effect on the computations of $\mathbf{A}^a = \mathbf{A}^f \mathbf{X}$ include all the volcanic-related ash/gas assimilation, e.g., assimilation of satellite data/LIDAR data/in situ data; (sand/desert) dust storm related assimilation; tornado-related assimilation; assimilation of exploding nuclear plants or factories; chemicals or oils leaking on seas; global (forecast) fire assimilation; assimilation of environmental pollutant transport, e.g., severe smog. In addition, for DA applications (e.g., Ozone, SO_2) where pollutants spread over the whole domain, usually the focus is only on the high

concentrations and a threshold can be set to ignore the very low values without losing the necessary assimilation accuracy. In this case, MS can also lead to a potential acceleration since many very low concentrations can be explicitly truncated to be zeros.

It has been analyzed that when the number of non-zero rows (n_{ash}, i.e., the number of ash grids in a 3-D domain) of $\mathbf{A}^f$ satisfies $n_{ash} < n$, MS can work faster than standard EnKF. For volcanic ash application, because n_{ash} is much less than n, the acceleration is thus quite large. Hence in this case, we propose to embed MS in all ensemble-based DA methods because it is fast and the implementation using MS is exact to the standard ensemble-based methods, i.e., it doesn't introduce any approximation in view of MS procedures. Actually this proposal can be extended to all real applications, even if the condition is not satisfied. This is because, in this case the computational cost of MS for $\mathbf{A}^a = \mathbf{A}^f \mathbf{X}$ becomes O($n N^2$), which is the same as that of using the standard assimilation (shown in Fig. 2(b)). Therefore, if the state numbers equal to or close to the number of the total number of grid points in the domain, the added computational cost by using MS is very small (neglectable), so that the computational time with MS is almost the same as the time of using the standard approach. Whereas, when the condition $n_{ash} < n$ is satisfied, MS will accelerate the analysis step. Thus MS is generic and can be directly used in any ensemble-based DA, and this acceleration can be automatically realized for some potential applications, without spending time investigating if the condition is satisfied. In a real (or operational) 3-D DA system, MS can be easily included, i.e., we only need to invoke the MS module when computing $\mathbf{A}^a = \mathbf{A}^f \mathbf{X}$, without any other change to the current framework.

As stated in Eq. (15), the speedup of the MS method is approximately the inverse of $\frac{n_{ash}}{n}$. So far there is no statistical data on the value of $\frac{n_{ash}}{n}$. Considering the problem of volcanic ash transport, there is one emission point (at the volcano), all the ashes in atmospheres are transported by the directional wind drive from the same source point. Thus volcanic ash cloud is actually transported in a shape of a plume, which in general doesn't cover the full but only a small part of the 3-D domain. At the start phase of a volcanic ash eruption, $\frac{n_{ash}}{n}$ is much smaller than 1.0 (started from 0). During transport over a long time (one and a half months for this case study), $\frac{n_{ash}}{n}$ increases to approximately $\frac{1}{3}$. Therefore, the speedup of MS on volcanic ash DA will be significant.

6.2 MS and localization

Based on the formulation of MS, one may think it can be taken as a localization approach ((Hamill et al., 2001)). There is indeed a similarity between MS and the localization approach, in a sense that when computing $\mathbf{A}^a = \mathbf{A}^f \mathbf{X}$, both get rid of a large number of cells, and only do computations related to the selected grids. These two algorithms are however functionally different. This is because the localization approach is meant for reducing spurious covariances outside a local region which is built up around the measurement, thus the results with and without localization approaches are different. While, MS is developed for the acceleration purpose. The masked region is discontinuous and independent of locations of measurement, but dependent on the model domain. Thus, there is no difference on the assimilation results between using MS and without using it. Therefore, based on the functional difference, MS cannot be taken as a localization approach.

In this study, we don't employ the localization strategy in the analysis step, because we use a rather large ensemble size of 100 to guarantee the accuracy, as introduced in Sect. 2. But for some applications (e.g., ozone, CO_2, sulfur dioxide) especially

when assimilating satellite data, localization is a necessary approach and has been widely used in reducing spurious covariances (Barbu et al., 2009; Chatterjee et al., 2012; Curier et al., 2012). In these cases, because the localization approach forces the analysis only to update state within a localization region, one may think that localization could replace MS and there would be no significance to employ MS. Actually this is not correct. We explain the reason as follows.

5   The localization approach is usually realized in Eq. (6) by employing a Schur product of a localization matrix and the forecast error covariance matrix (Houtekamer and Mitchell, 1998, 2001) given by:

$$\mathbf{K}(k) = (\mathbf{f} \circ \mathbf{P}^f(k))\mathbf{H}(k)^T[\mathbf{H}(k)(\mathbf{f} \circ \mathbf{P}^f(k))\mathbf{H}(k)^T + \mathbf{R}]^{-1}. \tag{17}$$

The Schur product $\mathbf{f} \circ \mathbf{P}^f$ in Eq. (17) is defined by the element-wise multiplication of the covariance matrix $\mathbf{P}^f$ and a localization matrix $\mathbf{f}$. $\mathbf{f}$ is defined based on the distance between two locations, thus it is dependent on the domain and needs
10   information of the full ensemble state locations. In this way, $\mathbf{f} \circ \mathbf{P}^f$ can contain more zeros than $\mathbf{P}^f$, but the dimensions are not changed, so that the computations related to $\mathbf{f} \circ \mathbf{P}^f$ are actually not reduced. Therefore, we can understand the localization approach in the analysis step as that the state within and outside a local region are both updated with increments, but just the increments outside the region are zero (which seems like not updating). This is also the reason why the localization approach is not meant for acceleration but only for reducing spurious covariances. Now it is clear that localization cannot replace MS.
15   Actually both can be performed together in dealing with the time-consuming part $\mathbf{A}^a = \mathbf{A}^f \mathbf{X}$. The localization approach can first transfer $\mathbf{A}^f$ to a localized matrix with more zero rows. Then MS can be used to accelerate the multiplication of the localized matrix and $\mathbf{X}$. In this way, MS is expected to accelerate $\mathbf{A}^a = \mathbf{A}^f \mathbf{X}$ with a high speedup rate, because the computational cost of more zero rows in the localized ensemble state matrix are avoided.

6.3   MS and parallelization

20   Motivated from the model's physics, the implementation MS currently is for the serial case. This implementation has reduced the computation time to an acceptable time (i.e., simulation time is less than the period of forecast in real world time). It is however interesting to discuss the potential of parallelization of the dense-dense matrix multiplication ($\tilde{\mathbf{A}}^a_{n_{ash} \times N} = \tilde{\mathbf{A}}^f_{n_{ash} \times N} \mathbf{X}_{N \times N}$) in the step (iv) of the algorithm (see Sect. 4.2 and Table 2). The related matrix multiplication can be easily parallelized on multiple processors. Optimization and evaluation on the parallelized MS will be considered in future. For the
25   current case study, the computational time (3.13 h, see Table 3) for an "ideal" reduction by parallelization of MS is not much larger than the acceleration (already) gained by MS (2.26 h, subtraction between 3.13 h and 0.87 h, see Table 3). Therefore, from application's perspective, further acceleration by parallelization is not required.

    Alternatively, one may also consider to (1) directly parallelize the expensive matrix multiplication of $\mathbf{A}^a_{n \times N} = \mathbf{A}^f_{n \times N} \mathbf{X}_{N \times N}$, without first performing MS; or (2) implement CSR-based SDMM (see Sect. 5) with parallelization. Both are possible alterna-
30   tive approaches to accelerate the expensive matrix multiplication. The first approach can be implemented by user own designed parallelization, or by utilizing scaLAPACK (https://www.netlib.org/scalapack/, where the main function is "pdgemm"). The second approach can be realized by using some general parallel sparse-dense matrix multiplication methods (e.g., sending each column of X and three CSR arrays of $\mathbf{A}^f$ to each processor to calculate each column of $\mathbf{A}^a$) or using a good parallel alge-

bra library like PeTSC (https://www.mcs.anl.gov/petsc/) which allows users to specify own orderings and comes with machine optimized parallel matrix-matrix multiplication operations. However, given the fact that MS can also be parallelized using similar ways or same libraries, thus it is fair to not consider parallelization for all cases (i.e., using MS, not using MS, using CSR-based SDMM). Actually, the parallelization in MS could be performed much easier than other approaches in dealing with $\mathbf{A}^a = \mathbf{A}^f \mathbf{X}$, because the dense-dense matrix multiplication (parallelization in the step (iv) of MS) is easier to parallelize than the sparse-dense matrix multiplication (direct parallelization for $\mathbf{A}^a = \mathbf{A}^f \mathbf{X}$ or parallelized CSR-based SDMM).

In this paper, for the current usage, we keep the possibility of parallelization open, because a serial MS has been efficient already.

7 Conclusions

In this study, based on evaluations on the computational cost of volcanic ash DA, the analysis step turned out to be very expensive. Although some potential approaches can accelerate the initialization and forecast step, there would be no notable improvement to the total computational time due to the dominant analysis step. Therefore, to get an acceptable computational cost, the key is to efficiently reduce the execution time of the analysis step.

After a detailed evaluation on various parts of the analysis stage, the most time-consuming part was revealed. The mask-state algorithm (MS) was developed based on a study on the characteristic of the ensemble ash states. The algorithm transforms the full ensemble state matrix into a relatively small matrix using a constructed mask array. Subsequently, the computation of the analysis step was sufficiently reduced. MS is developed as a generic approach, thus it can be embedded in all ensemble-based DA implementations. The extra computational cost of the algorithm is small and usually neglectable.

The conventional ensemble-based DA with MS is shown to successfully reduce the total computational time to an acceptable level, i.e., less than the time period of the assimilation application. Consequently, timely and accurate volcanic ash forecasts can be provided for aviation advices. This approach is flexible. It boosts the performance without considering any model-based parallelization, such as domain or component decomposition. Thus, when a parallel model is available, MS can be easily combined with the model to gain a further speedup. It implements exactly the standard DA without any approximation and with easy configurations, so that it can be used to accelerate the standard DA in a wide range of applications.

In this case study with the LOTOS-EUROS model (version 1.10), after the parallelization is performed for the forecast step of EnKF assimilation, the analysis step takes 72 % of total runtime, which means the analysis step is the bottle neck. This case might not be general for all ash forecasts. As the computational cost for initialization and forecast greatly depends on the forecast model that is used. For the current development, it makes sense to use the LOTOS-EUROS model, because the model has been configured and evaluated in (Fu et al., 2015) by comparison with other famous models (e.g., NAME (Jones et al., 2007), WRF-Chem (Webley et al., 2012)) in simulating volcanic ash transport. However, if a more expensive ash forecasting model is used, then the bottle neck would be the forecast step. In this case, the forecast step should be the goal for acceleration and probably a parallel model or adaptive domain-decomposition (as discussed in Sect. 4.3) needs to be employed together with the parallel ensemble forecasts.

The use of in situ measurements is one important reason why MS perfectly works. For each analysis step, the number of measurements are quite small, and the procedure of the singular value decomposition (SVD) costs little. However, in some applications when many measurements are assimilated (e.g., satellite-based or seismic-based data), and the number of measurements is of the same order as the number of state variables, the most time-consuming part will be the SVD. In these cases, the contributions of MS will be limited. The reduction of the total computing time using MS therefore is less significant, an effective acceleration algorithm for the analysis step must be used and should consider the computationally-expensive SVD in the first place.

8 Data and code availability

The averaged aircraft in situ data used in this study are available from Fig. 1(c). The used continuous aircraft data and the model output data can be accessed by request (G.Fu@tudelft.nl). The mask-state algorithm (MS) is implemented in OpenDA (the open source software for DA, www.openda.com) and the software can be downloaded from sourceforge (https://sourceforge.net/projects/openda).

17

5   *Competing interests.* The authors declare that they have no conflict of interest.

References

Amdahl, G. M.: Validity of the Single Processor Approach to Achieving Large Scale Computing Capabilities, in: Proceedings of the April 18-20, 1967, Spring Joint Computer Conference, AFIPS '67 (Spring), pp. 483–485, ACM, New York, NY, USA, doi:10.1145/1465482.1465560, 1967.

5 Bank, R. and Douglas, C.: Sparse matrix multiplication package (SMMP), 1, 127–137, doi:10.1007/bf02070824, 1993.

Barbu, A. L., Segers, A. J., Schaap, M., Heemink, A. W., and Builtjes, P. J. H.: A multi-component data assimilation experiment directed to sulphur dioxide and sulphate over Europe, Atmospheric Environment, 43, 1622–1631, doi:10.1016/j.atmosenv.2008.12.005, 2009.

Casadevall, T. J.: The 1989 – 1990 eruption of Redoubt Volcano, Alaska: impacts on aircraft operations, Journal of Volcanology and Geothermal Research, 62, 301–316, doi:10.1016/0377-0273(94)90038-8, 1994.

10 Chatterjee, A., Michalak, A. M., Anderson, J. L., Mueller, K. L., and Yadav, V.: Toward reliable ensemble Kalman filter estimates of CO2 fluxes, J. Geophys. Res., 117, D22 306+, doi:10.1029/2012jd018176, 2012.

Curier, R. L., Timmermans, R., Calabretta-Jongen, S., Eskes, H., Segers, A., Swart, D., and Schaap, M.: Improving ozone forecasts over Europe by synergistic use of the LOTOS-EUROS chemical transport model and in-situ measurements, Atmospheric Environment, 60, 217–226, doi:10.1016/j.atmosenv.2012.06.017, 2012.

15 Eliasson, J., Palsson, A., and Weber, K.: Monitoring ash clouds for aviation, Nature, 475, 455, doi:10.1038/475455b, 2011.

Evensen, G.: The Ensemble Kalman Filter: theoretical formulation and practical implementation, Ocean Dynamics, 53, 343–367, doi:10.1007/s10236-003-0036-9, 2003.

Filgueira, R., Atkinson, M., Tanimura, Y., and Kojima, I.: Applying Selectively Parallel I/O Compression to Parallel Storage Systems, in: Euro-Par 2014 Parallel Processing, edited by Silva, F., Dutra, I., and Santos Costa, V., vol. 8632 of *Lecture Notes in Computer Science*,

20 pp. 282–293, Springer International Publishing, doi:10.1007/978-3-319-09873-9_24, 2014.

Folch, A., Costa, A., Durant, A., and Macedonio, G.: A model for wet aggregation of ash particles in volcanic plumes and clouds: 2. Model application, J. Geophys. Res., 115, B09 202+, doi:10.1029/2009jb007176, 2010.

Fu, G., Lin, H. X., Heemink, A. W., Segers, A. J., Lu, S., and Palsson, T.: Assimilating aircraft-based measurements to improve Forecast Accuracy of Volcanic Ash Transport, Atmospheric Environment, 115, 170–184, doi:10.1016/j.atmosenv.2015.05.061, 2015.

25 Fu, G., Heemink, A., Lu, S., Segers, A., Weber, K., and Lin, H.-X.: Model-based aviation advice on distal volcanic ash clouds by assimilating aircraft in situ measurements, Atmospheric Chemistry and Physics, 16, 9189–9200, doi:10.5194/acp-16-9189-2016, 2016.

Gudmundsson, M. T., Thordarson, T., Höskuldsson, A., Larsen, G., Björnsson, H., Prata, F. J., Oddsson, B., Magnússon, E., Högnadóttir, T., Petersen, G. N., Hayward, C. L., Stevenson, J. A., and Jónsdóttir, I.: Ash generation and distribution from the April-May 2010 eruption of Eyjafjallajökull, Iceland, Scientific Reports, 2, doi:10.1038/srep00572, 2012.

30 Hamill, T. M., Whitaker, J. S., and Snyder, C.: Distance-Dependent Filtering of Background Error Covariance Estimates in an Ensemble Kalman Filter, Mon. Wea. Rev., 129, 2776–2790, doi:10.1175/1520-0493(2001)129%3C2776:ddfobe%3E2.0.co;2, 2001.

Houtekamer, P. L. and Mitchell, H. L.: Data Assimilation Using an Ensemble Kalman Filter Technique, Mon. Wea. Rev., 126, 796–811, doi:10.1175/1520-0493(1998)126%3C0796:dauaek%3E2.0.co;2, 1998.

Houtekamer, P. L. and Mitchell, H. L.: A Sequential Ensemble Kalman Filter for Atmospheric Data Assimilation, Mon. Wea. Rev., 129,

35 123–137, doi:10.1175/1520-0493(2001)129%3C0123:asekff%3E2.0.co;2, 2001.

Houtekamer, P. L., He, B., and Mitchell, H. L.: Parallel Implementation of an Ensemble Kalman Filter, Mon. Wea. Rev., 142, 1163–1182, doi:10.1175/mwr-d-13-00011.1, 2014.

Jones, A., Thomson, D., Hort, M., and Devenish, B.: The U.K. Met Office's Next-Generation Atmospheric Dispersion Model, NAME III, in: Air Pollution Modeling and Its Application XVII, edited by Borrego, C. and Norman, A.-L., pp. 580–589, Springer US, doi:10.1007/978-0-387-68854-1_62, 2007.

Kalnay, E., Ota, Y., Miyoshi, T., and Liu, J.: A simpler formulation of forecast sensitivity to observations: application to ensemble Kalman 5 filters, Tellus A, 64, doi:10.3402/tellusa.v64i0.18462, 2012.

Keppenne, C. L.: Data Assimilation into a Primitive-Equation Model with a Parallel Ensemble Kalman Filter, Mon. Wea. Rev., 128, 1971–1981, doi:10.1175/1520-0493(2000)128%3C1971:daiape%3E2.0.co;2, 2000.

Keppenne, C. L. and Rienecker, M. M.: Initial Testing of a Massively Parallel Ensemble Kalman Filter with the Poseidon Isopycnal Ocean General Circulation Model, Mon. Wea. Rev., 130, 2951–2965, doi:10.1175/1520-0493(2002)130%3C2951:itoamp%3E2.0.co;2, 2002.

10 Khairullah, M., Lin, H., Hanea, R. G., and Heemink, A. W.: Parallelization of Ensemble Kalman Filter (EnKF) for Oil Reservoirs with Time-lapse Seismic Data, International Journal of Mathematical, Computational Science and Engineering, 7, doi:waset.org/Publication/16317, 2013.

Liang, B., Sepehrnoori, K., and Delshad, M.: An Automatic History Matching Module with Distributed and Parallel Computing, Petroleum Science and Technology, 27, 1092–1108, doi:10.1080/10916460802455962, 2009.

15 Lin, H.-X., Cosman, A., Heemink, A., Stijnen, J., and van Beek, P.: Parallelization of the Particle Model SIMPAR, in: Advances in Hydro-Science and Engineering, edited by Holz, K. P., Bechteler, W., Wang, S. S. Y., and Kawahara, M., vol. 3, Center for Computational Hydro-science and Engineering, https://www.researchgate.net/publication/252671025_Parallelization_of_the_Particle_Model_SIMPAR, 1998.

Miyazaki, K., Eskes, H. J., and Sudo, K.: A tropospheric chemistry reanalysis for the years 2005–2012 based on an assimilation of OMI, MLS, TES, and MOPITT satellite data, Atmospheric Chemistry and Physics, 15, 8315–8348, doi:10.5194/acp-15-8315-2015, 2015.

20 Nerger, L. and Hiller, W.: Software for ensemble-based data assimilation systems—Implementation strategies and scalability, Computers & Geosciences, 55, 110–118, doi:10.1016/j.cageo.2012.03.026, 2013.

Oxford-Economics: The Economic Impacts of Air Travel Restrictions Due to Volcanic Ash, Report for Airbus, Tech. rep., http://www.oxfordeconomics.com/my-oxford/projects/129051, 2010.

Petrie, R. E. and Dance, S. L.: Ensemble-based data assimilation and the localisation problem, Weather, 65, 65–69, doi:10.1002/wea.505, 25 2010.

Quinn, J. C. and Abarbanel, H. D. I.: Data assimilation using a GPU accelerated path integral Monte Carlo approach, Journal of Computational Physics, 230, 8168–8178, doi:10.1016/j.jcp.2011.07.015, 2011.

Riishojgaard, L. P.: A direct way of specifying flow-dependent background error correlations for meteorological analysis systems, Tellus A, 50, 42–57, doi:10.1034/j.1600-0870.1998.00004.x, 1998.

30 Saad, Y.: Iterative Methods for Sparse Linear Systems, Society for Industrial and Applied Mathematics, doi:10.1137/1.9780898718003, 2003.

Schaap, M., Timmermans, R. M. A., Roemer, M., Boersen, G. A. C., Builtjes, P. J. H., Sauter, F. J., Velders, G. J. M., and Beck, J. P.: The LOTOS EUROS model: description, validation and latest developments, International Journal of Environment and Pollution, 32, 270+, doi:10.1504/ijep.2008.017106, 2008.

35 Segers, A. J.: Data Assimilation in Atmospheric Chemistry Models Using Kalman Filtering, Delft Univ Pr, http://repository.tudelft.nl/islandora/object/uuid%3A113b6229-c33a-4100-93be-22e1c8912672?collection=research, 2002.

Tavakoli, R., Pencheva, G., and Wheeler, M. F.: Multi-level Parallelization of Ensemble Kalman Filter for Reservoir History Matching, in: SPE Reservoir Simulation Symposium, Society of Petroleum Engineers, doi:10.2118/141657-ms, 2013.

Weber, K., Eliasson, J., Vogel, A., Fischer, C., Pohl, T., van Haren, G., Meier, M., Grobéty, B., and Dahmann, D.: Airborne in-situ investigations of the Eyjafjallajökull volcanic ash plume on Iceland and over north-western Germany with light aircrafts and optical particle counters, Atmospheric Environment, 48, 9–21, doi:10.1016/j.atmosenv.2011.10.030, 2012.

Webley, P. W., Steensen, T., Stuefer, M., Grell, G., Freitas, S., and Pavolonis, M.: Analyzing the Eyjafjallajökull 2010 eruption using satellite remote sensing, lidar and WRF-Chem dispersion and tracking model, J. Geophys. Res., 117, D00U26+, doi:10.1029/2011jd016817, 2012.

Zehner, C., ed.: Monitoring Volcanic Ash From Space, ESA communication Production Office, doi:10.5270/atmch-10-01, 2010.
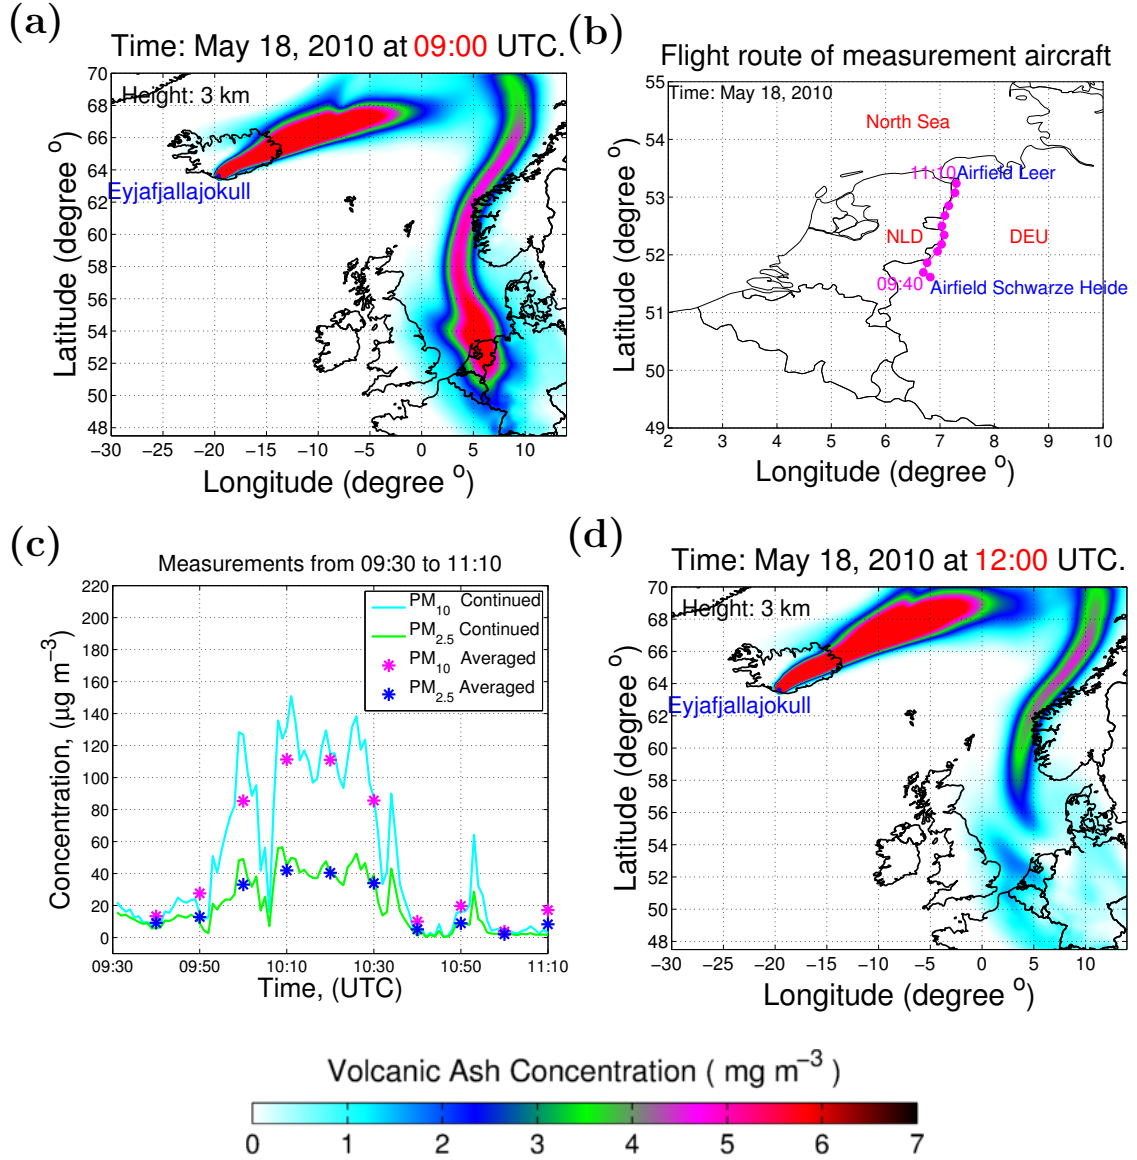
Figure 1. Methodology of ensemble-based DA. (a) The initial volcanic ash state at 09:00 UTC. **(b)** Flight route of measurement aircraft. **(c)** Aircraft in situ measurements of PM_{10} and $PM_{2.5}$ from 09:30 to 11:10 UTC May 18, 2010. **(d)** Volcanic ash assimilation result at 12:00 UTC.

(a)

n ash state number (3.888×10^6)
m measurement number (2)
N ensemble size (100)
$\mathbf{A}^f$ forecasted ensemble state matrix $(n \times N)$
$\mathbf{O}^f$ measured ensemble state matrix $(m \times N)$
$\mathbf{Y}$ measurement ensemble matrix $(m \times N)$
$\mathbf{A}^a$ analyzed ensemble state matrix $(n \times N)$

$\mathbf{B}_{N \times N}$ represents $\mathbf{I}_{N \times N} - \frac{1}{N}\mathbf{1}_{N \times N}$

Forecast step

$\mathbf{A}^f_{n \times N}$

Analysis step
$\mathbf{A}^a_{n \times N}$

$\mathbf{O}^f_{m \times N}$
$\mathbf{Y}_{m \times N}$

Stochastic Observer

(b)

Computational cost of analysis step

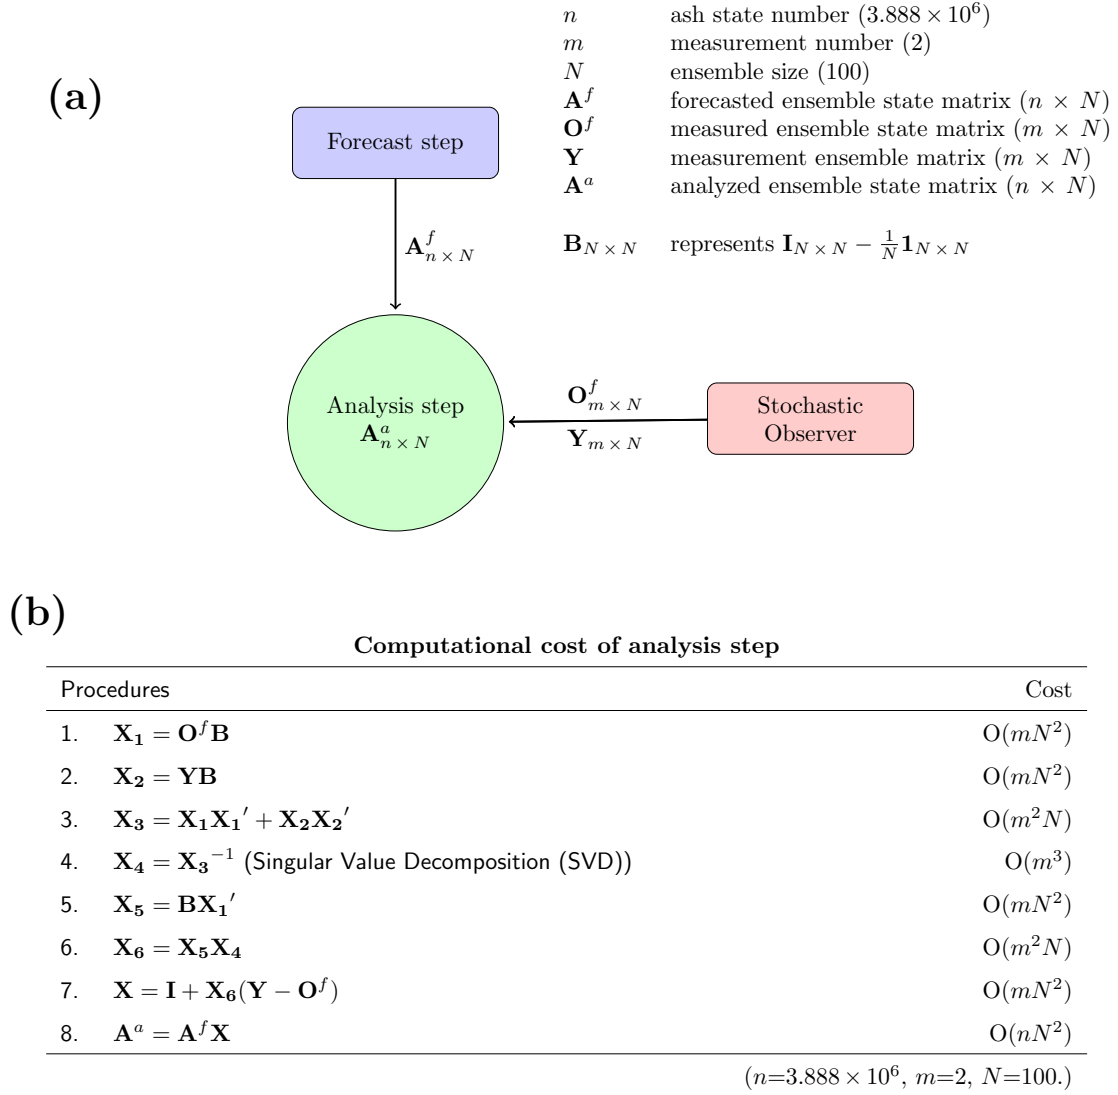| Procedures	Cost
1. $\mathbf{X_1} = \mathbf{O}^f \mathbf{B}$	$O(mN^2)$
2. $\mathbf{X_2} = \mathbf{YB}$	$O(mN^2)$
3. $\mathbf{X_3} = \mathbf{X_1}\mathbf{X_1}' + \mathbf{X_2}\mathbf{X_2}'$	$O(m^2 N)$
4. $\mathbf{X_4} = \mathbf{X_3}^{-1}$ (Singular Value Decomposition (SVD))	$O(m^3)$
5. $\mathbf{X_5} = \mathbf{B}\mathbf{X_1}'$	$O(mN^2)$
6. $\mathbf{X_6} = \mathbf{X_5}\mathbf{X_4}$	$O(m^2 N)$
7. $\mathbf{X} = \mathbf{I} + \mathbf{X_6}(\mathbf{Y} - \mathbf{O}^f)$	$O(mN^2)$
8. $\mathbf{A}^a = \mathbf{A}^f \mathbf{X}$	$O(nN^2)$

$(n=3.888 \times 10^6,\ m=2,\ N=100.)$

Figure 2. Computational evaluation of the analysis step. **(a)** Illustration of the analysis step. **(b)** Computational cost of all sub-part of the analysis step.

Figure 3. Characteristic of volcanic ash state.

(a)

<div align="center">CSR-based SpMV ($\mathbf{A}^f$ is represented in CSR format)</div>

Algorithm 1 $\quad \mathbf{A}^a[:,1]_{n \times 1}=$CSR-based SpMV$(\mathbf{A}^f_{n \times N}, \boldsymbol{v}_{N \times 1})$
1. for $i = 1 : n$
2. $\qquad \mathbf{A}^a[i,1] = 0$
3. $\qquad$ for $j = \boldsymbol{row\_ptr}[i] : \boldsymbol{row\_ptr}[i] - 1$
4. $\qquad\qquad w = \boldsymbol{col\_idx}[j]$
5. $\qquad\qquad \mathbf{A}^a[i,1] \mathrel{+}= \boldsymbol{val}[j] \times \boldsymbol{v}[w]$
6. $\qquad$ end for
7. end for

<div align="right">(n=3.888 $\times 10^6$, N=100.)</div>

(b)

<div align="center">SDMM is conducted by looping SpMV N times</div>

Algorithm 2 $\quad \mathbf{A}^a[:,1:N]_{n \times N} = N$-times **Algorithm 1**)
1. for $k = 1 : N$
2. $\qquad \mathbf{A}^a[:,k]_{n \times 1} = $ CSR-based SpMV$(\mathbf{A}^f_{n \times N}, \mathbf{X}[:,k]_{N \times 1})$
3. end for

<div align="right">(n=3.888 $\times 10^6$, N=100.)</div>

Figure 4. Algorithms for CSR-based SDMM to compute the multiplication of sparse matrix $\mathbf{A}^f$ and dense matrix X. (a) Multiplication of $\mathbf{A}^f$ with a column vector $\boldsymbol{v}$ (in **X**) by CSR-based sparse matrix vector multiplication (SpMV). $\boldsymbol{val}$, $\boldsymbol{col\_idx}$, $\boldsymbol{row\_ptr}$ are the three arrays to represent $\mathbf{A}^f$ in CSR format. **(b)** Looping SpMV N times (each with one column of **X**) to obtain $\mathbf{A}^a = \mathbf{A}^a \mathbf{X}$.

Table 1. Comparison of the computational cost of conventional EnKF and MS-EnKF. (The results are obtained from the bullx B720 thin nodes of the Cartesius cluster, which is a computing facility of SURFsara, the Netherlands Supercomputing Centre. Each node is configured with 2×12-core 2.6 GHz Intel Xeon E5-2690 v3 (Haswell) CPUs and with memory 64 GB.)

Case	Conventional EnKF	MS-EnKF
Cores used	102	102
Tracer number (n_{spec})	6	6
Measurements of tracers (m)	2	2
Ensemble size (N)	100	100
Parallel in forecast step	Yes	Yes
Parallel in analysis step	No	No
Mask-state in analysis step	**No**	**Yes**
Initialization	0.42 h	0.42 h
Forecast	0.65 h	0.65 h
Analysis	**3.14** h	**0.88** h
Others	0.15 h	0.12 h
Total Runtime	**4.36** h	**1.95** h

h=hour, simulation window = **3.0** h, the time is Wall clock time.

Table 2. Computational evaluation of all the steps of the mask-state algorithm (MS) for $\mathbf{A}^a = \mathbf{A}^f \mathbf{X}$. (see details of each step of MS in Sect. 4.2.)

Sub-step	Computational time
(i) Compute ensemble mean state $\bar{\mathbf{A}}^f$	0.0097 h
(ii) Construct mask array $\mathbf{z}$	0.0002 h
(iii) Construct masked ensemble state matrix $\tilde{\mathbf{A}}^f$	0.0057 h
(iv) Compute $\tilde{\mathbf{A}}^a$ by multiplying $\tilde{\mathbf{A}}^f$ and $\mathbf{X}$	**0.8474** h
(v) Construct analyzed ensemble state matrix $\mathbf{A}^a$	0.0070 h
Total	**0.87** h

h=hour, the time is Wall clock time.

Table 3. Computational time for the analysis step of conventional EnKF, MS-EnKF, CSR-based-SDMM-EnKF.

Analysis's procedures (see Fig. 2(b))	Conventional EnKF	MS-EnKF	CSR-based-SDMM-EnKF
procedures 1 - 8	0.01 h	0.01 h	0.01 h
procedure 9 ($\mathbf{A}^a = \mathbf{A}^f \mathbf{X}$)	**3.13** h	**0.87** h	**1.21** h
Total	**3.14** h	**0.88** h	**1.22** h

h=hour, the time is Wall clock time.

Table 4. Computational evaluation of the sub-steps of the Sparse-Dense Matrix Multiplication with Compressed Sparse Row storage (CSR-based SDMM) for $\mathbf{A}^a = \mathbf{A}^f \mathbf{X}$.

Sub-step	Computational time
(i) Compute three arrays (in CSR format) of $\mathbf{A}^f$	0.0407 h
(ii) Compute CSR-based SpMV for the first column of $\mathbf{A}^a$	0.0117 h
(iii) Loop (ii) for N-1 times for other N-1 columns of $\mathbf{A}^a$	**1.1576** h
Total	**1.21** h

CSR-based SDMM is formed by (ii) and (iii). h=hour, the time is Wall clock time.