



A mechanism for catastrophic filter divergence in data assimilation for sparse observation networks

G. A. Gottwald¹ and A. J. Majda²

¹School of Mathematics and Statistics, University of Sydney, NSW 2006, Australia

²Department of Mathematics and Center for Atmosphere Ocean Science, Courant Institute of Mathematical Sciences, New York University, USA

Correspondence to: G. A. Gottwald (georg.gottwald@sydney.edu.au)

Received: 28 March 2013 – Revised: 31 July 2013 – Accepted: 13 August 2013 – Published: 25 September 2013

Abstract. We study catastrophic filter divergence in data assimilation procedures whereby the forecast model develops severe numerical instabilities leading to a blow-up of the solution. Catastrophic filter divergence can occur in sparse observational grids with small observational noise for intermediate observation intervals and finite ensemble sizes. Using a minimal five-dimensional model, we establish that catastrophic filter divergence is a numerical instability of the underlying forecast model caused by the filtering procedure producing analyses which are not consistent with the true dynamics, and stiffness caused by the fast attraction of the inconsistent analyses towards the attractor during the forecast step.

1 Introduction

Data assimilation is the procedure to find the best estimation of the state of a dynamical system given a forecast model with possible model error and noisy observations at discrete observation intervals (Kalnay, 2002; Majda and Harlim, 2012). The presence of the often chaotic nature of the underlying nonlinear dynamics, as well as the sparseness of the observational network, significantly complicates this process. In the setting of ensemble-based filters (Evensen, 1994, 2006), finite ensemble sizes may introduce additional sources of error (see, for example, Ehrendorfer, 2007). Insufficient ensemble size typically causes an underestimation of the error variances, which may ultimately lead to filter divergence when the filter trusts its own forecast and ignores the information provided by the observations. This filter divergence is caused by ensemble members aligning with the

most unstable directions (Ng et al., 2011) and is exacerbated by large observational noise. Finite size effects may also lead to spurious overestimating correlations between otherwise uncorrelated variables (Hamill et al., 2001; Whitaker et al., 2004, 2009; Liu et al., 2008; Sacher and Bartello, 2008), spoiling the overall analysis skill.

Harlim and Majda (2010) and Gottwald et al. (2011) documented a new type of filter divergence which is characterised by the forecast model diverging to machine infinity. It was shown that this catastrophic filter divergence occurs in sparse observational networks with small observational noise for moderate observation intervals, in contrast to the classical filter divergence described in the previous paragraph.

We will establish here the mechanism leading to this instability in a minimal low-dimensional model: in a sparse observational grid, finite ensemble sizes cause the ensemble to align, and in the case of (sufficiently) small observational noise generate analyses which are not consistent with the actual dynamics and are located in phase space off the attractor. If the attraction towards the attractor is sufficiently strong, the subsequent forecast step attempts to integrate a stiff dynamical system, which may cause the integrator to develop numerical instabilities.

In Sect. 2 we introduce the minimal model for which catastrophic filter divergence is studied. We briefly describe ensemble Kalman filters in Sect. 3. Numerical results are presented in Sect. 4, and the mechanism for catastrophic filter divergence is established. We conclude with a discussion in Sect. 5.

2 A minimal model

We study the Lorenz-96 (Lorenz, 1996) model

$$\dot{z}_i = z_{i-1}(z_{i+1} - z_{i-2}) - z_i + F \quad i = 1, \dots, D \quad (1)$$

with $\mathbf{z} = (z_1, \dots, z_D)$ and periodic $z_{i+D} = z_i$ in a five-dimensional setting. We use negative forcing here, which allows strong mixing with small dimension D . We consider here $D = 5$ with $F = -16$ (Abramov and Majda, 2006). For these parameters we find as Lyapunov exponents $\lambda = (2.72, 0.09, -0.09, -1.83, -5.89)$ for an integration lasting 250 time units. One of the Lyapunov exponents should be zero, corresponding to the flow direction; due to slow convergence this is only approximately satisfied. Note that $\sum_{i=1}^5 \lambda_i = \lim_{t \rightarrow \infty} \frac{1}{t} \int \text{Tr}(M(t)) dt$, where M is the linearized vector field of Eq. (1), and hence $\sum_{i=1}^5 \lambda_i = -5$. Using the Kaplan–Yorke dimension (see for example Schuster and Just, 2005), this suggests that the attractor has a fractal dimension of $D_{\text{attr}} = 4.15$, and trajectories are on average attracted to this manifold with the fast rate $\lambda_5 = -5.89$. The climatic mean and variance are estimated from a long time trajectory as $\bar{z} = -2.47$ and $\sigma_{\text{clim}}^2 = 33.7$, respectively.

We remark that the system in Eq. (1) has not been chosen to model any physical system but rather for its simplicity in addressing the phenomenon of catastrophic blow-up. We will report as well on results with $F = 8$, which is less chaotic with $\sigma_{\text{clim}}^2 = 13.1$, $\lambda = (0.474, 0.003, -0.523, -1.315, -3.636)$ and $D_{\text{attr}} = 2.9$.

We assume that observations of the variables are given at equally spaced discrete observation times t_i with observation interval Δt_{obs} . We observe only one variable z_1 . It is well known that the Kalman filter is suboptimal for dynamical systems which are nonlinear and involve non-Gaussian statistics. It is pertinent to mention that although the five-dimensional Lorenz system in Eq. (1) is highly nonlinear, its probability density function is near-Gaussian for $F = -16$, but highly non-Gaussian for $F = 8$. The Lorenz system in Eq. (1) is assimilated using an ensemble transform Kalman filter (ETKF) (Tippett et al., 2003; Wang et al., 2004), which is briefly described in the following section.

3 Ensemble Kalman filter

In an ensemble Kalman filter (EnKF) (Evensen, 2006), an ensemble with k members $\mathbf{z}_k \in \mathbb{R}^D$

$$\mathbf{Z} = [\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_k] \in \mathbb{R}^{D \times k}$$

is propagated by the full nonlinear dynamics

$$\dot{\mathbf{Z}} = \mathbf{F}(\mathbf{Z}), \quad \mathbf{F}(\mathbf{Z}) = [f(z_1), f(z_2), \dots, f(z_k)] \in \mathbb{R}^{D \times k}. \quad (2)$$

The ensemble is split into its mean

$$\bar{\mathbf{z}} = \frac{1}{k} \sum_{i=1}^k \mathbf{z}_i,$$

and its ensemble deviation matrix

$$\mathbf{Z}' = \mathbf{Z} - \bar{\mathbf{z}} \mathbf{e}^T,$$

where $\mathbf{e} = [1, \dots, 1]^T \in \mathbb{R}^k$. The ensemble deviation matrix \mathbf{Z}' is used to provide a Monte Carlo estimate of the forecast covariance matrix

$$\mathbf{P}_f(t) = \frac{1}{k-1} \mathbf{Z}'(t) \mathbf{Z}'(t)^T \in \mathbb{R}^{D \times D}.$$

Note that $\mathbf{P}_f(t)$ is rank-deficient if the ensemble size k is smaller than the rank of the covariance matrix.

Given the forecast ensemble $\mathbf{Z}_f = \mathbf{Z}(t_i - \epsilon)$ and the associated forecast error covariance matrix (or the prior) $\mathbf{P}_f(t_i - \epsilon)$, the actual Kalman analysis (Kalnay, 2002; Evensen, 2006; Simon, 2006) updates a forecast into a so-called analysis (or the posterior). Variables at times $t = t_i - \epsilon$ are evaluated before taking observations \mathbf{y}_o into account in the analysis step, and variables at times $t = t_i + \epsilon$ are evaluated after the analysis step when the observations, taken at $t = t_i$, have been taken into account. Observations $\mathbf{y}_o \in \mathbb{R}^n$ can be expressed as a perturbed truth according to

$$\mathbf{y}_o(t_i) = \mathbf{H}\mathbf{z}(t_i) + \mathbf{r}_o,$$

where the observation operator $\mathbf{H} : \mathbb{R}^D \rightarrow \mathbb{R}^n$ maps from the whole space into observation space, and $\mathbf{r}_o \in \mathbb{R}^n$ is i.i.d. observational Gaussian noise with associated error covariance matrix \mathbf{R}_o and zero mean.

In the first step of the analysis, the forecast mean $\bar{\mathbf{z}}_f$ is updated to the analysis mean

$$\bar{\mathbf{z}}_a = \bar{\mathbf{z}}_f - \mathbf{K}_o [\mathbf{H}\bar{\mathbf{z}}_f - \mathbf{y}_o], \quad (3)$$

where the Kalman gain matrix is defined as

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

The analysis covariance \mathbf{P}_a is given by

$$\mathbf{P}_a = (\mathbf{I} - \mathbf{K}_o \mathbf{H}) \mathbf{P}_f. \quad (5)$$

To calculate an ensemble \mathbf{Z}_a which is consistent with the analysis error covariance \mathbf{P}_a , and which therefore needs to satisfy

$$\mathbf{P}_a = \frac{1}{k-1} \mathbf{Z}_a \mathbf{Z}_a^T,$$

we use the method of deterministic ensemble square root filters (Simon, 2006) which expresses the analysis ensemble as a linear combination of the forecast ensemble. In particular we use the method proposed in Tippett et al. (2003) and Wang et al. (2004), the so-called ensemble transform Kalman filter (ETKF). Alternatively one could have chosen the ensemble adjustment filter (Anderson, 2001) or the continuous Kalman–Bucy filter, which does not require the inversion of matrix inverses (Bergemann et al., 2009). A new forecast $\mathbf{Z}(t_{i+1} - \epsilon)$ is then obtained by propagating \mathbf{Z}_a with the full nonlinear dynamics to the next time of observation. The numerical results presented in the next section are obtained with this method.

4 The genesis of catastrophic filter divergence

We observe only one of the five variables z_i (without loss of generality we use z_1) and generate observations y_o from the truth by adding Gaussian observational noise with small observational error covariance $\mathbf{R}_o = 0.01$ after equal observation intervals Δt_{obs} . We use in the following a fourth-order Runge–Kutta method to integrate forward in time the system in Eq. (1) during the forecasting step.

In Fig. 1 we show an instance of catastrophic filter divergence for $dt = 0.025$ and $\Delta t_{\text{obs}} = 0.05$ where we used $k = 6$ ensemble members (so the forecast error covariance matrix is not necessarily rank-deficient). Besides the maximal absolute amplitude of the analysis ensemble, we show the norm error \mathcal{E} of the analysis

$$\mathcal{E}(t_i) = \|\bar{z}_a(t_i) - z_t(t_i)\| \tag{6}$$

evaluated at each analysis cycle t_i between the truth z_t and the ensemble mean \bar{z}_a . After $t_i = 62$ the norm error becomes machine infinity, due to the forecast model developing a numerical instability. The genesis of the blow-up is clearly seen from Fig. 1: until $t_1 \approx 55.5$ the filter is stable and the analysis tracks where the norm error may be even smaller than the observational error (see the inset in Fig. 1). This is followed by a non-tracking episode lasting to $t_i \approx 59$ in which the norm error evolves around a mean value of approximately $\langle \mathcal{E} \rangle \approx 18 \approx \sqrt{\langle \mathcal{E}^2 \rangle - \text{Var}[\mathcal{E}]} = \sqrt{2D\sigma_{\text{clim}}^2 - \text{Var}[\mathcal{E}]}$, suggesting that the analysis is exploring the attractor, uncorrelated from the truth and not controlled by the observations anymore. This episode precedes the actual blow-up episode of the forecast integrator in which the norm error grows to machine infinity.

In order to quantify the propensity for blow-up, we count the number N_b of blow-ups that occur before a total of 5000 simulations have terminated without blow-up. A single successful simulation consists of $n_a = 4000$ analysis cycles. Simulations differ in truth, observations and in the initial ensemble with variance 1. The proportions of blow-ups is then given by

$$S_b = \frac{N_b}{N_b + 5000}.$$

Note that S_b depends on the number of data assimilation cycles n_a . In Fig. 2 we show S_b as a function of the observation interval Δt_{obs} for several values of the integration time step dt . It is seen that blow-up occurs for moderate observation time intervals. No blow-up occurs for sufficiently small or sufficiently large values of Δt_{obs} . The percentage of blow-ups as well as the range of Δt_{obs} for which blow-up occurs is reduced by reducing the integration time step, establishing blow-up as a numerical instability of the forecast model. Additionally, we performed simulations with the forcing parameter in Eq. (1) chosen as $F = 8$, corresponding to less chaotic dynamics. We found similar behaviour; however, in




Fig. 1. Top: maximal absolute value Z_{amax} of the analysis ensemble over all $D = 5$ components for $dt = 0.025$ and $\Delta t_{\text{obs}} = 0.05$ with $\mathbf{R}_o = 0.01$ and $k = 6$ ensemble members. Bottom: the error norm \mathcal{E} as a function of analyses cycles. The continuous line (online blue) in the inset shows the observational error $\sqrt{\mathbf{R}_o}$.

line with the less chaotic nature of the system when compared to $F = -16$, blow-up develops for larger values of the integration time step dt . All results presented in the following are for $F = -16$.

To obtain meaningful statistics for blow-up which are independent of the number of the analyses cycles n_a , we estimate the number of assimilation cycles τ_i before blow-up occurs. To generate statistics of these blow-up times, we numerically calculated τ_i for 10 000 instances of blow-up where we allowed for a maximum of $n_a = 250\,000$ assimilation cycles. In Fig. 3 we show the empirical cumulative distribution function $P_b(\tau_i)$ for the blow-up times. The results suggest that catastrophic blow-up is a Poisson process with cumulative probability distribution function




Fig. 2. Percentage S_b of blow-ups as a function of the observation interval Δt_{obs} for several values of the integration time step dt for fixed number of analyses cycles $n_a = 4000$. We used $\mathbf{R}_o = 0.01$ and $k = 6$ ensemble members.

$$P_b(\tau_i) = 1 - \exp\left(-\frac{\tau_i}{\bar{\tau}_b}\right), \quad (7)$$

where $\bar{\tau}_b$ denotes the mean blow-up time. The Poisson character of blow-up suggests that blow-up is not an accumulative process, but rather that blow-up is a random process with each assimilation having the same probability of blow-up, independent of previous assimilations. Linear regression of the curves in Fig. 3 yields for the average blow-up times $\bar{\tau}_b = 62633$ ($\Delta t_{\text{obs}} = 0.005$), $\bar{\tau}_b = 3324$ ($\Delta t_{\text{obs}} = 0.05$) and $\bar{\tau}_b = 6160$ ($\Delta t_{\text{obs}} = 0.1$). This is consistent with the blow-up statistics S_b for fixed number of assimilation cycles n_a and shows that the probability of blow-up has a maximum for an intermediate value of the observation interval Δt_{obs} for a specified integration time step dt (cf. red curve with circles in Fig. 2).

We now look at the dependency of the propensity for blow-up on the observational variance \mathbf{R}_o . Figure 4 shows that catastrophic filter blow-up requires the observational noise to be sufficiently small but non-zero. For the smallest value of the observational noise we used with $\mathbf{R}_o = 10^{-6}$, we still observed $S_b = 0.07$ for $n_a = 4000$, $dt = 0.0025$ and $\Delta t_{\text{obs}} = 0.05$. This is in stark contrast to the traditional filter divergence which occurs for sufficiently large observational noise.

We propose that catastrophic filter divergence is caused by insufficient ensemble size paired with sufficiently small observational noise. We have checked that by increasing the ensemble size to impractically high values of $k = 400$, we were able to avoid catastrophic blow-up. To monitor the ensemble spread, we consider the ensemble dimension D_{ens} as defined in Patil et al. (2001); Pazó et al. (2011),




Fig. 3. Log plot of the empirical normalised cumulative probability density function $P_b(\tau_i)$ of blow-up times τ_i for different observation intervals Δt_{obs} for $dt = 0.0025$. We used $\mathbf{R}_o = 0.01$ and $k = 6$ ensemble members.

$$D_{\text{ens}} = \frac{\left(\sum_{i=1}^k \sqrt{\mu_i}\right)^2}{\sum_{i=1}^k \mu_i},$$

where μ_i denotes the i th eigenvalue of the $k \times k$ covariance matrix

$$C = X_f^T X_f.$$

Note that D_{ens} takes values between 1 and $\min(k, D)$, depending on whether the ensemble members are all aligned or are orthogonal to each other. In Fig. 5 we show the ensemble dimension as a function of time for an ensemble with $k = 6$ members corresponding to the blow-up presented in Fig. 1. It is seen that $D_{\text{ens}} \approx 2$ during the stable tracking episode, indicating that the ensemble is not spanning all directions on the attractor (we recall the fractal attractor dimension to be $D_{\text{attr}} = 4.15$) but instead is aligning with the first two Lyapunov vectors (cf. Ng et al., 2011). This triggers the non-tracking period until $t \approx 59$. On the other hand, for ensemble sizes of $k = 400$ we observe that mostly $D_{\text{ens}} > 4$, and no blow-up occurs. Before the actual blow-up the ensemble dimension reaches values of almost $D_{\text{ens}} = 1$, indicating ensemble collapse.

Finite ensemble sizes and the associated loss of ensemble spread are known to cause non-catastrophic filter divergence in which the filter trusts the wrong forecasts, ignoring error-correcting observations (Houtekamer and Mitchell, 1998; Hamill et al., 2001; Sacher and Bartello, 2008; Ng et al., 2011). Finite ensemble sizes cause the forecast error covariance \mathbf{P}_f to exhibit on the one hand small diagonal variances and on the other hand off-diagonal entry values

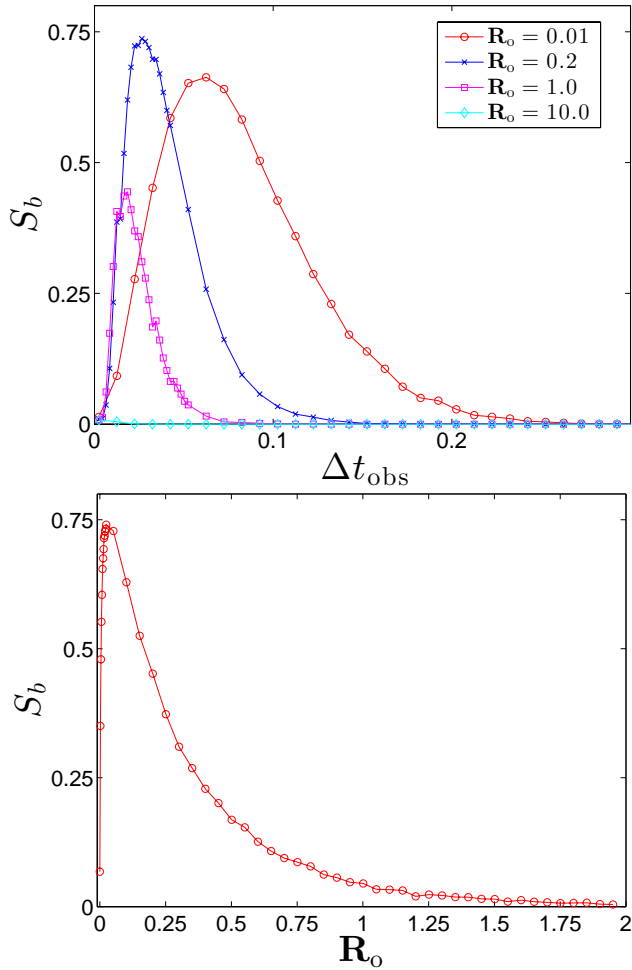


Fig. 4. Top: percentage S_b of blow-ups as a function of the observation interval Δt_{obs} for several values of the observational noise \mathbf{R}_o for fixed number of analyses cycles $n_a = 4000$, integration time step $dt = 0.0025$. Bottom: percentage S_b of blow-ups as a function of the observational noise variance \mathbf{R}_o for fixed number of analyses cycles $n_a = 4000$, integration time step $dt = 0.0025$ and observational interval $\Delta t_{\text{obs}} = 0.05$.

of unrealistically large absolute value (Hamill et al., 2001). This leads to unrealistic innovations of the unobserved variables towards the observation of the observed distant variable. Gottwald et al. (2011) showed that catastrophic filter divergencies are suppressed by a variance limiting Kalman filter (VLKF) which controls overestimation of the analysis error covariance.

The destructive interplay of sparse (sufficiently) accurate observations and finite size ensembles can be illustrated as follows. The Kalman filter produces analyses according to Eq. (3), which read for our case where only z_1 is observed as

$$\bar{z}_{ai} = \bar{z}_{fi} - \frac{\mathbf{P}_{f1i}}{\mathbf{P}_{f11} + \mathbf{R}_o} [\bar{z}_{f1} - y_o], \quad (8)$$




Fig. 5. Ensemble dimension D_{ens} as a function of analyses cycles. Parameters as in Fig. 1.

for $i = 1, \dots, 5$. The combination of small ensemble sizes causing small values of \mathbf{P}_{f11} and large absolute values of \mathbf{P}_{f1i} for $i > 1$ and comparably small observational noise \mathbf{R}_o leads to analyses for $i \neq 1$ which are significantly influenced by the observation y_o at site $i = 1$, irrespective of the actual physical correlations present in the dynamics. The resulting analysis may therefore not be dynamically consistent but may lie in phase space off the attractor. As a proxy for the distance of the analysis to the attractor, we measure the time τ_{attr} taken for the trajectory to reach an Euclidean distance θ from the attractor when propagating the analysis forward in time. We created an approximation of the attractor by storing 2×10^6 data points sampled at 0.005 time units. We choose $\theta = 1$. Figure 6 illustrates clearly that during the non-tracking period episode (i.e. the period $t_i \in (55, 59)$ in Fig. 1), one has (predominantly) $\tau_{\text{attr}} = 0$, consistent with our previous observation that the analysis lies on the attractor, but is uncorrelated from the truth and not controlled by the observations. The subsequent initiation of blow-up for $t > 59$, however, is characterised by non-zero values of τ_{attr} . It is clearly seen that blow-up is characterised by analyses lying off the attractor. It is pertinent to mention that the existence of alignment of the ensemble and the occurrence of off-attractor analyses (i.e. large values of τ_{attr}) does not necessarily cause catastrophic filter divergence (for example as in Fig. 6 at $t_i \approx 56$). The effect of off-attractor analyses is the following: the forecast model, initialised with such an analysis lying off the attractor, tries to follow the stable direction towards the globally attracting set with a rate which is in our case very fast on average with a Lyapunov exponent of -5.89 . This renders the dynamical system stiff developing numerical instabilities for sufficiently large time steps dt , causing the filter to catastrophically diverge to machine infinity.




Fig. 6. Distance between analysis and the attractor as measured by τ_{attr} . Parameters as in Fig. 1.

As seen in Fig. 2 there is no filter divergence for sufficiently small and sufficiently large observation intervals Δt_{obs} . This can now be understood as follows: for too small observation intervals, the forecast model will not have sufficiently propagated the analysis away from dynamically realistic values, whereas for sufficiently large values of $\Delta t_{obs} \gg \tau_{corr}$ the ensemble will have acquired sufficient spread with $D_{ens} \geq 4$ exploring the whole of the attractor.

The behaviour of the propensity S_b for blow-up as a function of the noise covariance \mathbf{R}_o as depicted in Fig. 4b can also be readily understood from Eq. (8); for large observational noise with $\mathbf{R}_o > 2$ we obtain analyses $\bar{z}_a = \bar{z}_f$ which are forced by the dynamics to lie on the attractor, and the sampling error of finite ensemble sizes is not entering the analysis. For small observational noise the magnitude of the increments $(\mathbf{P}_{f1i}/(\mathbf{P}_{f11} + \mathbf{R}_o)) [\bar{z}_{f1} - y_o]$ is given as a balance between small innovations $[\bar{z}_{f1} - y_o]$ and large finite ensemble size induced gains $\mathbf{P}_{f1i}/(\mathbf{P}_{f11} + \mathbf{R}_o) \approx \mathbf{P}_{f1i}/\mathbf{P}_{f11}$, yielding a maximum at $\mathbf{R}_o \approx 0.025$.

We now address the question how catastrophic blow-up can be controlled, except through decreasing the time step of the numerical forecast model to control the numerical instability or through increasing the number of ensembles to diminish the sampling effects. We found that catastrophic blow-up can be avoided by employing covariance localisation into the data assimilation procedure, which controls the unrealistic overestimation of off-diagonal entries of the forecast covariance matrix \mathbf{P}_f . Houtekamer and Mitchell (2001) and Hamill et al. (2001) achieved covariance localisation by Schur multiplication of the forecast error covariance \mathbf{P}_f with a localisation matrix C_{loc} . We used the compactly supported localisation function introduced by Gaspari and Cohn (1999), in conjunction with a DEnKF proposed by Sakov and Oke (2008), and found that catastrophic filter diver-




Fig. 7. Percentage S_b of blow-ups as a function of the multiplicative inflation factor δ for fixed number of analyses cycles $n_a = 4000$, integration time step $dt = 0.0025$ and observational noise variance $\mathbf{R}_o = 0.01$. We used $k = 6$ ensemble members.

gence is completely suppressed for a localisation radius of $\rho_{loc} < 1.5$ grid spacings. Other covariance limiting strategies such as the ensemble filters suggested in Gottwald et al. (2011) and Mitchell and Gottwald (2012) were also able to suppress catastrophic filter divergence. We remark that the actual truth, however, does indeed exhibit nontrivial correlations between all variables for our parameters in this low dimension with $D = 5$. Multiplicative error covariance (i.e. the multiplication of the forecast error covariance matrix \mathbf{P}_f by a constant inflation factor δ) is a standard remedy to control underestimation of covariances (Anderson and Anderson, 1999). In Fig. 7 we show the effect of multiplicative covariance inflation on the propensity for blow-up. We observe that increasing the inflation factor δ from $\delta = 1$ significantly decreases the propensity for blow-up S_b ; this is achieved by reducing occurrences of non-tracking periods (the classical filter divergence), which are the precursors of catastrophic blow-up. For sufficiently large values of δ , however, the instances of catastrophic filter divergence are drastically increased. Too large values of the inflation factor exacerbate the sampling errors of the forecast error covariances with $\delta \mathbf{P}_{f1i}/(\delta \mathbf{P}_{f11} + \mathbf{R}_o) \approx \mathbf{P}_{f1i}/\mathbf{P}_{f11}$. Covariance inflation, however, cannot completely suppress catastrophic filter divergence for all observation times, and more complicated behaviour can occur as observed for example for $\Delta t_{obs} = 0.04$ in Fig. 7.

5 Discussion

We have numerically established that catastrophic filter divergence is caused by the interplay of finite size effects and sparse observations with small observational noise producing analyses which may be situated in phase space far away from the actual attractor. The subsequent attraction back towards the attractor by the forecast model may cause numerical instabilities if the attraction rate is sufficiently large. This suggests that blow-up is to be expected in sparse observational networks involving observables which exhibit a large degree of irregularity. If those high variance fields are measured sufficiently accurately, catastrophic filter divergence is possible. This is, for example, the case in data assimilation of small-scale intermittent turbulent fields or in situations where sparse accurate observations of variables exhibiting strong spatial gradients such as jets can cause numerical instabilities to occur (J. L. Anderson, personal communication, 2012).

We have checked that our results are independent of the numerical integration scheme used during the forecasting step. We have performed simulations with a first-order in time forward Euler scheme and a second-order in time implicit midpoint rule scheme. The latter is unconditionally stable for the system in Eq. (1) (see Theorem 5.5.6. in Stuart and Humphries, 1996); however, in the case of the implicit midpoint rule, we observed a large increase of the iterations required to solve the nonlinear fixed point equation, rendering the scheme impractical. Furthermore, we have performed simulations with the deterministic ensemble adjustment Kalman filter (EAKF) (Anderson and Anderson, 1999) and observed similar behaviour.

Our work established the dynamical genesis of catastrophic filter divergence. Besides an impractical reduction of the integration time step (or an increase of the limit of iterations required in an unconditionally stable implicit method), to control the stiffness of the dynamical system, or an impractical increase of the number of ensembles to eliminate finite size sampling errors, covariance localisation and appropriately tuned covariance inflation were found to be effective in mitigating catastrophic filter divergence. Our study further shows that choosing too small observation covariances can lead to filter blow-ups as observed in numerical weather prediction models (A. Shlyayeva, personal communication, 2012).

Acknowledgements. We are grateful to Jeff Anderson, Heikki Jarvinen, Balu Nadiga and Anna Shlyayeva for sharing their experiences with catastrophic filter divergences in their numerical simulations. We are grateful for the detailed referee report by Jeff Anderson, which pointed us towards a possible Poisson statistics of blow-ups. G. A. Gottwald acknowledges support from the Australian Research Council, and A. J. Majda from ONR grants N0014-11-1-0306, DRI grant, N00014-10-1-0554, and MURI grant, N00014-12-1-0912.

Edited by: O. Talagrand

Reviewed by: J. Anderson and one anonymous referee

References

- Abramov, R. V. and Majda, A. J.: New Approximations and Tests of Linear Fluctuation-Response for Chaotic Nonlinear Forced-Dissipative Dynamical Systems. *J. Nonlin. Sci.*, 18, 303–341, 2006.
- Anderson, J. L.: An ensemble adjustment Kalman filter for data assimilation. *Mon. Weather Rev.*, 129, 2884–2903, 2001.
- Anderson, J. L. and Anderson, S. L.: A Monte Carlo implementation of the nonlinear filtering problem to produce ensemble assimilations and forecasts, *Mon. Weather Rev.*, 127, 2741–2758, 1999.
- Bergemann, K., Gottwald, G. A., and Reich, S.: Ensemble propagation and continuous matrix factorization algorithms. *Q. J. Roy. Meteorol. Soc.*, 135, 1560–1572, 2009.
- Ehrendorfer, M.: A review of issues in ensemble-based Kalman filtering, *Meteorol. Z.*, 16, 795–818, 2007.
- Evensen, G.: Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics, *J. Geophys. Res.*, 99, 10143–10162, 1994.
- Evensen, G.: *Data Assimilation: The Ensemble Kalman Filter*. Springer, New York, 2006.
- Gaspari, G. and Cohn, S. E.: Construction of correlation functions in two and three dimensions. *Q. J. Roy. Meteorol. Soc.*, 125, 723–757, 1999.
- Gottwald, G. A., Mitchell, L., and Reich, S.: Controlling overestimation of error covariance in ensemble Kalman filters with sparse observations: A variance limiting Kalman filter, *Mon. Weather Rev.*, 139, 2650–2667, 2011.
- Hamill, T. M., Whitaker, J. S., and Snyder, C.: Distance-dependent filtering of background covariance estimates in an ensemble Kalman filter, *Mon. Weather Rev.*, 129, 2776–2790, 2001.
- Harlim, J. and Majda, A. J.: Catastrophic filter divergence in filtering nonlinear dissipative systems, *Communications in Mathematical Sciences*, 8, 27–43, 2010.
- Houtekamer, P. L. and Mitchell, H. L.: Data assimilation using an ensemble Kalman filter technique. *Mon. Weather Rev.*, 126, 796–811, 1998.
- Houtekamer, P. L. and Mitchell, H. L.: A sequential ensemble Kalman filter for atmospheric data assimilation, *Mon. Weather Rev.*, 129, 123–136, 2001.
- Kalnay, E.: *Atmospheric Modeling, Data Assimilation and Predictability*, Cambridge University Press, Cambridge, 2001.
- Keating, S. R., Majda, A. J., and Smith, K. S.: New methods for estimating ocean eddy heat transport using satellite altimetry, *Mon. Weather Rev.*, 140, 1703–1722, 2012.
- Liu, J., Fertig, E. J., Li, H., Kalnay, E., Hunt, B. R., Kostelich, E. J., Szunyogh, I., and Todling, R.: Comparison between Local Ensemble Transform Kalman Filter and PSAS in the NASA finite volume GCM – perfect model experiments, *Nonlin. Processes Geophys.*, 15, 645–659, doi:10.5194/npg-15-645-2008, 2008.
- Lorenz, E. N.: *Predictability – a problem partly solved*, edited by: Palmer, T., in: *Predictability*, European Centre for Medium-Range Weather Forecast, Shinfield Park, Reading, UK, 1996.
- Majda, A. J. and Harlim, J.: *Filtering Complex Turbulent Systems*, Cambridge University Press, Cambridge, 2012.

- Mitchell, L. and Gottwald, G. A.: Controlling model error of underdamped forecast models in sparse observational networks using a variance-limiting kalman filter, *Q. J. Roy. Meteorol. Soc.*, 139, 212–225, 2012.
- Ng, G.-H. C., McLaughlin, D., Entekhabi, D., and Ahanin, A.: The role of model dynamics in ensemble Kalman filter performance for chaotic systems, *Tellus A*, 63, 958–977, 2011.
- Patil, D. J., Hunt, B., Kalnay, E., Yorke, J., and Ott, E.: Local low dimensionality of atmospheric dynamics. *Phys. Rev. Lett.*, 86, 5878–5881, 2001.
- Pazó, D., Rodríguez, M. A., and López, J. M.: Maximizing the statistical diversity of an ensemble of bred vectors by using the geometric norm, *J. Atmos. Sci.*, 68, 1507–1512, 2011.
- Sacher, W. and Bartello, P.: Sampling errors in ensemble Kalman filtering. Part I: Theory, *Mon. Weather Rev.*, 136, 3035–3049, 2008.
- Sakov, P. and Oke, P.: A deterministic formulation of the ensemble Kalman filter: An alternative to ensemble square root filters, *Tellus A*, 60, 415–428, 2008.
- Schuster, H. G. and Just, W.: *Deterministic Chaos: An Introduction*. John Wiley & Sons, 4th edition, 2005.
- Simon, D. J.: *Optimal State Estimation*, John Wiley & Sons, Inc., New York, 2006.
- Stuart, A. M. and Humphries, A.: *Dynamical Systems and Numerical Analysis*, Cambridge University Press, Cambridge, 1996.
- Tippett, M. K., Anderson, J. L., Bishop, C. H., Hamill, T. M., and Whitaker, J. S.: Ensemble square root filters, *Mon. Weather Rev.*, 131, 1485–1490, 2003.
- Wang, X., Bishop, C. H., and Julier, S. J.: Which is better, an ensemble of positive-negative pairs or a centered spherical simplex ensemble?, *Mon. Weather Rev.*, 132, 1590–1505, 2004.
- Whitaker, J. S., Compo, G. P., Wei, X., and Hamill, T. M.: Reanalysis without radiosondes using ensemble data assimilation. *Mon. Weather Rev.*, 132, 1190–1200, 2004.
- Whitaker, J. S., Compo, G. P., and Thépaut, J. N.: A comparison of variational and ensemble-based data assimilation systems for reanalysis of sparse observations. *Mon. Weather Rev.*, 137, 1991–1999, 2009.