



Data assimilation in the laboratory using a rotating annulus experiment

R. M. B. Young* and P. L. Read

Atmospheric, Oceanic and Planetary Physics, Department of Physics, University of Oxford, UK

*Correspondence to: AOPP, Clarendon Laboratory, Parks Road, Oxford OX1 3PU, UK. E-mail: young@atm.ox.ac.uk.

The thermally-driven rotating annulus is a laboratory experiment important for the study of the dynamics of planetary atmospheres under controllable and reproducible conditions. We use the analysis correction method to assimilate laboratory data into an annulus model. We analyse the 2S and 3AV regular flow regimes between rotation rates of 0.75 and 0.875 rad s⁻¹, and the 3SV chaotic flow regime between rotation rates of 2.2 and 3.1 rad s⁻¹. Our assimilated observations are irregularly-distributed, which is more meteorologically realistic than gridded observations, as used in recent applications of data assimilation to laboratory measurements. We demonstrate that data assimilation can be used successfully and accurately in this context. We examine a number of specific assimilation scenarios: a wavenumber transition between two regimes, information propagation from data-rich to data-poor regions, the response of the assimilation to a strong disturbance to the flow, and a vortex shedding instability phenomenon at high rotation rate. At the highest rotation rates we calculated the barotropic E-vectors using unobserved variables such as temperature and the vertical structure of the velocity field that are only available via the assimilation. These showed that the mean flow is weakened by the action of eddies, going some way towards explaining why vortices are shed at the very highest rotation rates but not at lower rotation. Rossby wave stability theory suggests the underlying instability leading to vortex shedding may be baroclinic in character.

Key Words: rotating annulus; analysis correction; data assimilation; E-vector; chaos; baroclinic instability

Received 27 April 2012; Revised 5 September 2012; Accepted 26 September 2012.

This is the peer reviewed version of the following article: Young RMB, Read PL. 2013. Data assimilation in the laboratory using a rotating annulus experiment. Q. J. R. Meteorol. Soc., 139, 1488–1504. DOI:10.1002/qj.2061, which has been published in final form at <http://dx.doi.org/10.1002/qj.2061>. This article may be used for non-commercial purposes in accordance with Wiley Terms and Conditions for Self-Archiving.

1. Introduction

The *analysis* is a meteorological term referring to the best estimate of a system's true state. It cannot be known exactly in practice because observations have finite density and intrinsic measurement error, but it can be estimated by combining sources of information about a system in a statistically and dynamically robust fashion; this is *data assimilation*. Data assimilation combines these sources in a way that minimizes the combined error, usually in a least-squares or maximum-likelihood sense. In the atmospheric sciences these sources are usually observations, a numerical model, climate distributions, and theoretical relationships between different physical quantities.

Data assimilation is most commonly used to produce initial conditions for weather and climate forecasts. It is also used to estimate the state in regions with few or no observations by using the model to pass information from data-rich to data-poor areas (Thompson 1961), to retrieve unobserved variables, identify weaknesses in the forecast model and the observational network, produce surrogates of the true state of the system for further study, and for *reanalysis*, where old data covering an extended period are assimilated into a model using a single method.

Daley (1991), Kalnay (2003), and others provide good overviews of the underlying principles and details of many different assimilation methods. The earliest methods simply interpolated observations onto a model grid (Panofsky 1949; Gilchrist and Cressman 1954). *Sequential* methods correct a first guess from climate or a previous forecast using observations as they become available over time, with the correction from each observation depending on the relative errors between a first guess and that observation

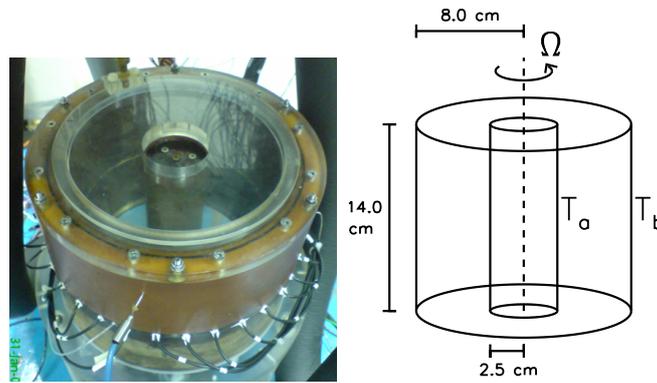


Figure 1. Left: Photograph of a rotating annulus experiment used in the AOPP fluid dynamics laboratory (image courtesy A. A. Castrejón-Pita). Right: Schematic of the experiment. Inner and outer cylinders at radii $R = a, b$ are at temperatures T_a and T_b respectively, rotating at constant angular velocity Ω , and fluid is contained between the cylinders.

(Bergthórsson and Döös 1955; Barnes 1964; Lorenc *et al.* 1991). Variational methods minimize a cost function based on the relative sizes of observation and first guess errors with respect to the system state (Sasaki 1970; Lorenc *et al.* 2000; Rawlins *et al.* 2007). Ensemble methods exploit the statistics of a forecast ensemble (Evensen 1994; van Leeuwen 2010). Finally, gradient descent filters look for model trajectories that ‘shadow’ observations using the full nonlinear forecast model (Stemler and Judd 2009; Judd *et al.* 2008). Data assimilation is a vibrant area of research, and there is as much need now for new approaches to these problems as there was in the 1960s.

The rotating annulus (Fig. 1) is a laboratory experiment that has been used for several decades to study the fundamental mechanisms underlying weather and climate under controlled and reproducible laboratory conditions. In its ‘standard’ configuration, a fluid annulus is contained between two coaxial circular cylinders. One heats the outer cylinder and cools the inner cylinder, while rotating the apparatus about its central vertical axis. This mimics the main influences acting on an Earth-like planetary atmosphere: gravity, rotation of the planet, and a temperature gradient between low and high latitudes. The experiment reproduces a wide range of Hadley cell-type, Rossby wave-type, vacillating, chaotic, and turbulent flow regimes depending on the fluid, dimensions, and forcing parameters (Hide and Mason 1975; Hignett *et al.* 1985; Read *et al.* 1992). In general the flow reflects the large-scale atmospheric flow patterns observed in the mid-latitudes of the Earth (and other terrestrial planets; Fultz 1951, for example). As the processes governing the atmospheric circulation and the laboratory rotating tank are similar, then similar methods may be used to study their behaviour. Applying such methods in one context may provide insights into their use in the other.

In this paper we use data assimilation to assimilate a number of rotating annulus observational datasets into a numerical model. This is a necessary part of a larger body of work studying the predictability of the rotating annulus experiment, which will be reported on elsewhere. Here we apply data assimilation in isolation to study the experiment and explore phenomena using variables only accessible via the assimilation process. For example, our observations contain a vortex shedding phenomenon at high rotation rates, whose behaviour cannot be understood without unobserved variables such as the fluid’s vertical structure and temperature. We use the observations to demonstrate that data can be assimilated accurately into a model of the rotating annulus, measuring the accuracy of the assimilation against an independent set of velocity observations. While being able to do this with an appropriately-tuned method is not too surprising with hindsight, it is not immediately clear *a priori* that this is straightforward given that annulus observations are generally highly incomplete.

We use the *analysis correction* method (Lorenc *et al.* 1991, hereafter L91), which is a well-established sequential assimilation technique. The UK Met Office used it operationally from 30 November 1988 (L91, p.82) until implementing 3D-Var on 29 March 1999 (Lorenc *et al.* 2000). For inclusion within a larger framework focusing on predictability rather than the properties of the method itself, the method is attractive for three main reasons. First, it is a relatively simple assimilation technique whose implementation is considerably cheaper than more advanced techniques, for only a small loss in accuracy. Second, groups at Oxford and the Open University use it in the context of Martian atmospheric data assimilation (Lewis *et al.* 1996; Montabone *et al.* 2006), so there is some local expertise on its use. Finally, a linearized version of our rotating annulus model was not available when the algorithm was designed, which precluded use of variational methods like 4D-Var.

Using our algorithm we produce sequences of analyses over three 3-hour sets of laboratory observations in two regular (non-chaotic) and one chaotic flow regime. The main features of the regular flow regimes (wavenumber-2 steady flow 2S and wavenumber-3 amplitude vacillation 3AV) are a baroclinic wave and an amplitude vacillation cycle over a period of several hundred seconds (Hide and Mason 1975, Fig. 5, for example). We also analyse a chaotic flow regime (wavenumber-3 structural vacillation 3SV), which is characterised by a baroclinic wave whose shape changes over time but whose amplitude remains approximately constant.

We know of two previous attempts at data assimilation in rotating laboratory experiments. First, Galmiche *et al.* (2003) and Thivolle-Cazat *et al.* (2005) studied vortex evolution at the CORIOLIS facility, a 13 m diameter open rotating tank in Grenoble. They used an extended Kalman filter method on their experimental data to assimilate it into an ocean model. Second, Ravela *et al.* (2010) assimilate rotating annulus data into a stripped-down General Circulation Model using the ensemble Kalman filter (Evensen 1994). Their experimental and model setup is similar to ours but with a few important differences. They use a hydrostatic model while ours is Boussinesq non-hydrostatic, and their model grid is regularly spaced while we use a stretched grid to resolve the boundary layers. Their work is impressive as they can assimilate observations and update the model in real time before recording the next observations, but updating the model in real time is not a priority in our work.

The most important difference between this work and the two previous attempts is the method used to obtain the observations. Galmiche *et al.* (2003) used correlation image velocimetry and Ravela *et al.* (2010) used particle image velocimetry to obtain

Table 1. Summary of the annulus, fluid, and model properties used by MORALS in this work. The setup is identical to the ‘main comparison’ of [Hignett et al. \(1985\)](#). The SI contains expressions showing how ν , κ , and ρ vary with T .

Annulus properties	
Inner cylinder radius	$a = 2.5$ cm
Outer cylinder radius	$b = 8.0$ cm
Depth of fluid	$d = 14.0$ cm
Model properties	
Timestep	$\delta t = 0.02$ s (expf2) $\delta t = 0.01$ s (expf5, expf6)
Grid points (radial, azimuthal, vertical)	$(N_R, N_\phi, N_z) = (24, 64, 24)$
Independent variables	121,088
Fluid properties (all at 20 °C)	
Working fluid (b.v.)	83% water, 17% glycerol
Density	$\rho = 1.044$ g cm ⁻³
Kinematic viscosity	$\nu = 1.715 \times 10^{-2}$ cm ² s ⁻¹
Thermal diffusivity	$\kappa = 1.284 \times 10^{-3}$ cm ² s ⁻¹
Volume expansion coefficient	$\alpha_v = 2.755 \times 10^{-4}$ degC ⁻¹
Prandtl number	$Pr = 13.4$

79 gridded velocity observations. Presenting the observations to the assimilation in grid form is not meteorologically realistic, because in
80 operational practice the observation network is not regularly spaced and many of the observation points move with time. In this work
81 we use observations both irregularly spaced and changing in position over time.

82 A final aim of this work, although this particular paper does not address this directly, is to set the foundations for a broader goal
83 in which the annulus is used as a testbed for the study and evaluation of new and proposed meteorological techniques not yet in
84 operational use. The controllable and reproducible conditions the laboratory annulus provides allow such methods to be tested under
85 more rigorous conditions than can be provided solely using atmospheric data. The controllable complexity is particularly important
86 because as the complexity of a system is reduced any imperfections in methods used to study it become less obscured by the complexity
87 of the system itself. Many studies of meteorological techniques use very low-dimensional systems like the [Lorenz \(1963\)](#) equations,
88 and the laboratory scale is a natural bridge between these low-order models and the atmosphere itself, where methods can be tested
89 using a real fluid with a non-idealised model and incomplete and noisy observations.

90 In the next section we describe our rotating annulus model and the laboratory data to be assimilated. Section 3 describes our
91 algorithm for data assimilation using analysis correction. In Sect. 4 we present a series of assimilations using these observations.
92 In Sect. 5 we examine the vortex shedding phenomenon at the highest rotation rates, and in Sect. 6 we conclude. The Supporting
93 Information (SI) contains more technical details about various aspects of the work, in particular the algorithm itself.

94 2. Model and observations

95 The numerical model used to simulate the rotating annulus experiment is the *Met Office / Oxford Rotating Annulus Laboratory*
96 *Simulation* (MORALS) ([Farnell and Plumb 1976](#)). The model is well established as a quantitatively accurate model of annulus flow in
97 regular and weakly chaotic flow regimes.

98 MORALS solves the Navier-Stokes, mass continuity, and heat transfer equations for a Boussinesq fluid rotating at angular velocity Ω
99 in a cylindrical fluid annulus, along with a Poisson equation for pressure, a quadratic equation of state for density, and two ‘constitutive’
100 relations for viscosity and thermal diffusivity. Apart from the constitutive relations there is no sub-grid scale parameterization. The
101 integration scheme is finite-difference accurate to $O(\delta t^2)$, where δt is the model time step, using a Leapfrog scheme with Robert-
102 Asselin filter. The equations are discretised on an Arakawa-C grid using cylindrical polar coordinates (R, ϕ, z) . The grid is non-uniform,
103 stretched in R and z to resolve the boundary layers.

104 The equations are cast in velocity-temperature-pressure form. There are four prognostic variables: u (radial), v (azimuthal) and w
105 (vertical) velocities / cm s⁻¹, and temperature T / °C. A fifth field, kinetic pressure $\Pi \equiv p/\rho_0$ / cm² s⁻², is diagnostic, calculated from
106 the other four fields using a Poisson equation. The setup is the ‘standard’ thermally-driven annulus configuration (Fig. 1) with inner
107 and outer walls maintained at constant temperatures T_a and T_b respectively (temperature difference $\Delta T = T_b - T_a$), constant rotation
108 rate Ω , and no internal heating. All velocities are set to zero at the boundaries (rigid lid), and the temperature gradient is zero across
109 the top and bottom boundaries (insulating). Table 1 lists the parameters defining the annulus and working fluid, which are identical to
110 [Hignett et al.’s \(1985\)](#) ‘main comparison’, and the numerical model.

111 We took laboratory data for this work from the AOPP archive of rotating annulus experimental data from 1998, using the expf2,
112 expf5, and expf6 *datafiles*, each of which contains horizontal velocity observations (radial and azimuthal) grouped in *datasets* every 5 s
113 (Fig. 2). The SI contains some details about the data pre-processing.

114 The apparatus used allowed measurements at up to five vertical levels in the annulus, 12.4, 9.7, 7.0, 4.3 and 1.6 cm above the bottom
115 of the tank. [Jackson and Hignett \(1984\)](#) and [Dalziel \(1995\)](#) detail the method used to take observations. A horizontal plane of light
116 ~ 5 mm thick shines into the tank, and reflects off neutrally buoyant particles into a video camera 1 m directly above the tank, rotating
117 in the same coordinate frame and about the same axis. The camera tracks the particles (100–500 per frame) over one second, producing
118 a set of images from which particle positions and velocities are calculated. The sequence of observations is staggered: data at the

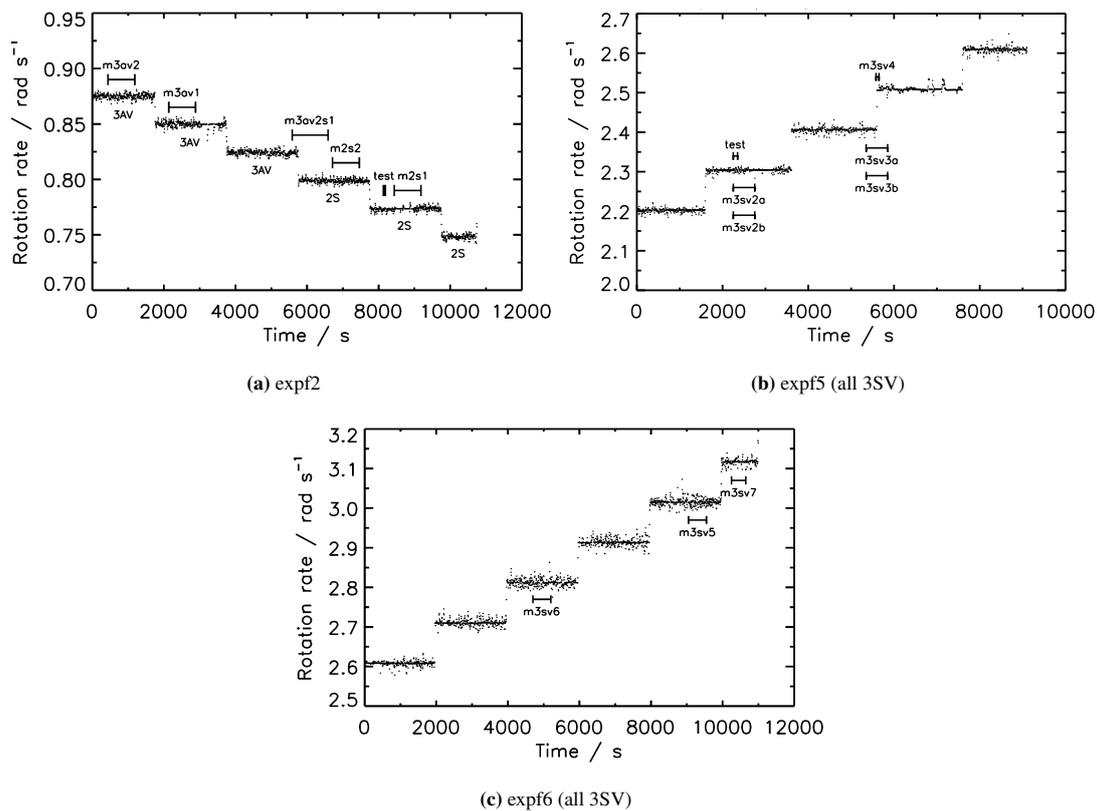


Figure 2. Rotation rate Ω as a function of time during the three datafiles. The plots also show the position of each assimilation and the parameter tests listed in Table 2 and presented in Sect. 4, and the flow regime at each rotation rate. Note that the scatter probably represents measurement error and not real scatter in the forcing.

119 top level are taken during the first 5 s, followed by 5 s at each of the other heights in turn, then the cycle repeats. No observations of
 120 vertical velocity, temperature, or pressure are available. Each velocity observation has an error associated with it, which is estimated
 121 to be 0.0057 cm s^{-1} in each of the radial and azimuthal directions (see SI for details of this estimate). This error is much smaller
 122 than the variability in the observed velocities, which is roughly 0.06 cm s^{-1} for the radial velocities and 0.09 cm s^{-1} for the azimuthal
 123 velocities.

124 Each dataset is split into two or more *subsets*, separated by approximately one second. While some individual measurements may
 125 be the same particle measured twice a second apart, the subsets are sufficiently uncorrelated to be used as independent data. For the
 126 purposes of the assimilation we assume the subsets are valid at the same time, which simplifies the algorithm and interpretation of
 127 results. The approximation is justified because the shortest dynamical timescale is significantly longer than a second.

128 expf2 starts as wavenumber-3 amplitude vacillation (3AV regime) and finishes about 11 000 s later as a wavenumber-2 steady wave
 129 (2S). It uses data at all five levels (Fig. 3), so consecutive datasets at a particular level are separated by 25 s. ΔT is approximately
 130 constant at around 4.05 degC while Ω decreases step-wise from 0.875 to 0.75 rad s^{-1} (Fig. 2a). There are two subsets within this data,
 131 ‘blue’ and ‘red’ (the SI contains the precise times for each one within the 5 s window).

132 expf5 and expf6 both contain about 10 000 s of horizontal velocity data in the chaotic wavenumber-3 structural vacillation (3SV)
 133 flow regime, Ω stepping up as shown in Fig. 2. ΔT is about 4.02 degC throughout. Only two levels are observed in these datafiles, at
 134 $z = 9.7 \text{ cm}$ and 4.3 cm (Fig. 4), so consecutive datasets at each level are separated by 10 s. There are three subsets in expf5 (‘blue’,
 135 ‘red’, and ‘orange’), and two in expf6 (‘blue’ and ‘red’).

136 One interesting feature of the expf5 and expf6 datafiles is that the neutrally buoyant particles used to track the flow velocity moved
 137 so slowly within the cyclones that they dropped out of suspension there (see SI Fig. 7). The piles formed on the base of the tank then
 138 caused the main baroclinic wave to become phase-locked to the tank, possibly by stimulating a topographic Rossby wave secondary to
 139 the main wave, with this feedback leading to more particles being deposited in the same place. These piles are not accounted for by the
 140 free-running model, of course.

141 As is clear from Figs 3 and 4, there are more observations at the top of the tank than at the bottom (approximately a factor of four
 142 difference). This is because the data collection software rejects more particle tracks at the lower levels. The light reflecting off the
 143 particles and into the camera has to pass through more liquid the lower in the tank the light sheet is, so more particle tracks at lower
 144 levels will be obscured by particles higher in the tank. The variation over time in the number of observations in each dataset is roughly
 145 5–10%.

146 In each subset roughly 500 data points are sampled, measuring both radial and azimuthal velocities. The ratio of observations to
 147 grid points is comparable with atmospheric models: compare $1000/121088 \sim 1\%$ in this context with $10^5/10^7 \sim 1\%$ for a typical
 148 atmospheric model (Kalnay 2003, p. 13), and note, again, that some quantities are not observed at all.

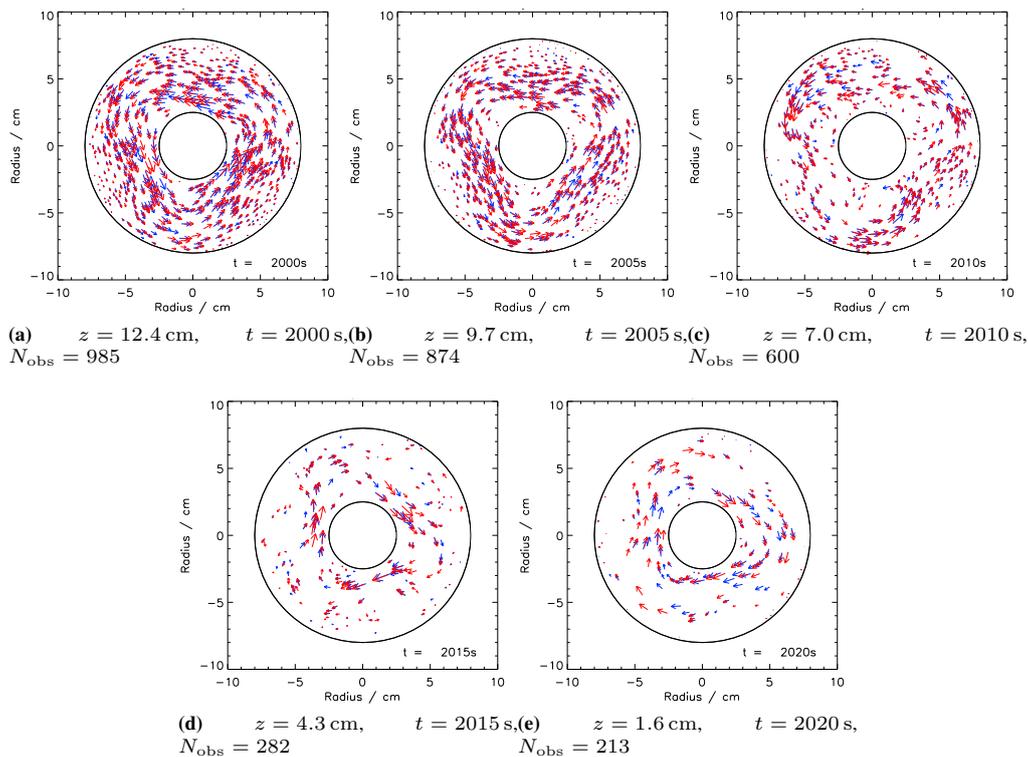


Figure 3. Representative horizontal velocity fields from datafile expf2 at 0.85 rad s^{-1} (3AV), also showing the number of observations N_{obs} at each vertical level. The blue vectors (black in print) show the observations from the subset at 2.8 s through the 5 s window, and the red vectors (grey in print) show the observations at 3.6 s.

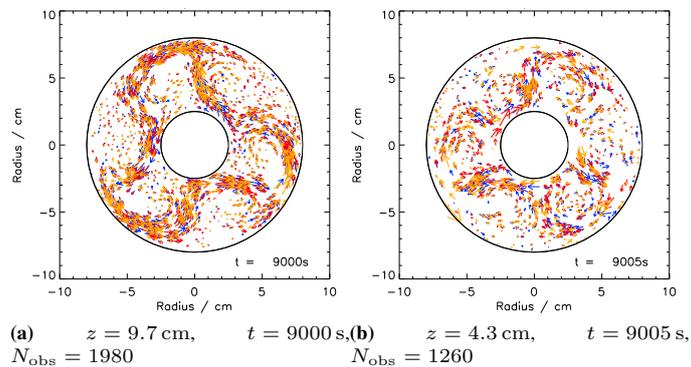


Figure 4. Representative horizontal velocity fields from datafile expf5 at 2.6 rad s^{-1} . The blue vectors (dark grey in print) show the observations at 0.7 s through the 5 s window, the red vectors (black in print) at 1.5 s, and the orange vectors (light grey in print) at 2.3 s. N_{obs} is the number of observations in each figure.

149 3. Assimilation Code for Analysis Correction In the Annulus

150 In this section we describe the method used to perform data assimilation on our annulus data. The core of this method is the analysis
 151 correction (AC) method developed in L91. The analysis computed by AC, denoted \mathbf{x}^a , represents the best estimate of the true state of
 152 the annulus at the *analysis time* t^a given information from observations, short range forecasts, and dynamical relationships between the
 153 different physical quantities. Each assimilation is done using the *Assimilation Code for Analysis Correction In the Annulus* (ACACIA).
 154 L91 contains more details on AC's underlying principles, and the SI contains full details of its application to MORALS.

155 In analysis correction, we step through the observational datafile in sequence. At each analysis time, the method takes a first guess
 156 or *background state* \mathbf{x}^b , compares it with the observed values at the observation points \mathbf{y}^o , and maps the difference between the two
 157 onto the model grid using various weighting functions, to give the analysis state \mathbf{x}^a . The model then runs forward for a short time using
 158 the analysis as the initial condition, up to the next analysis time, when the model run's end state is used as the next \mathbf{x}^b .

159 The original background state is taken from a free-running MORALS run, with the fields rotated in azimuth to line up with the
 160 observations. The observations used in each assimilation are taken from a window stretching forward in time by t^b and backwards
 161 in time by t^f relative to the analysis time (or t^b back in time and t^f forward in time relative to a particular observation), and each
 162 observation's contribution is weighted depending on how far the observation is from each grid point, and how close it is in time to the
 163 analysis time.

164 The assimilation window is longer than the time between analyses, so individual observations are used multiple times. This is called
 165 'forward continuous repeated insertion' (Fig. 5). Each observation enters the assimilation window with small influence over a large

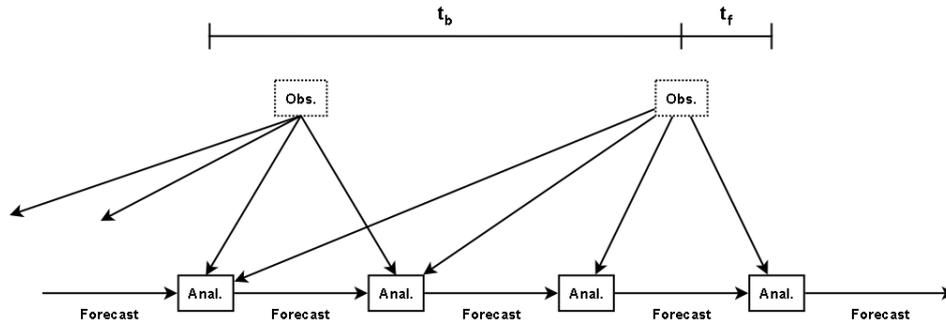


Figure 5. Schematic showing the elements of the analysis cycle, repeated insertion, and the observations contributing to each analysis. The arrows represent the flow of information, and time increases to the right. t^b and t^f are the times forward and backwards during which an observation influences the assimilation.

166 area (Fig. 6), and as the analysis approaches the observation time, the observation exerts more influence over a smaller area. When the
 167 analysis time has passed the observation time its influence falls away again until it leaves the window. Over time, the analysis converges
 168 to a state that is dynamically consistent with the sequence of observations both forward and backwards in time.

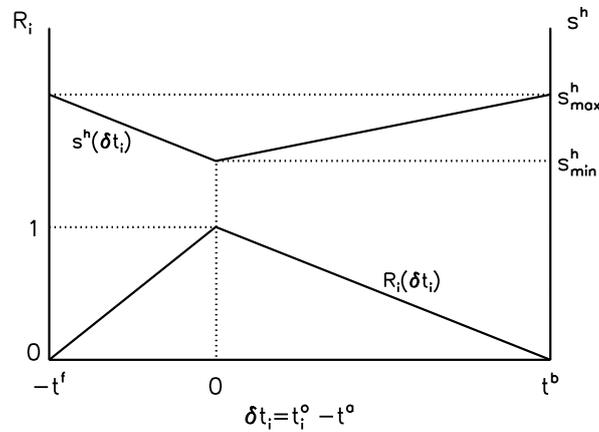


Figure 6. Temporal weighting factor $R_i(\delta t_i)$ and horizontal correlation scale s^h (similarly for vertical scale s^v) for each observation in the assimilation window, relative to the analysis time.

169 Each application of the method produces an analysis according to

$$\mathbf{x}^a = \mathbf{x}^b + \mathbf{W}\tilde{\mathbf{Q}} \left\{ \mathbf{y}^o - H(\mathbf{x}^b) \right\} \quad (1)$$

170 where H interpolates from the model grid to the observation points, and the gain matrix $\mathbf{W}\tilde{\mathbf{Q}}$ is approximated in component form by
 171 (L91, Eqs A1.9 and 3.21)

$$[\mathbf{W}\tilde{\mathbf{Q}}]_{ki} = \frac{\mu_{ki}}{\varepsilon_i^2 + (1 + \varepsilon_i^2)^{1/2} D_i} \equiv \mu_{ki} \tilde{Q}_i \quad (2)$$

172 where μ_{ki} is a combination of horizontal and vertical spatial weighting functions between model grid point k and observation i , D_i is
 173 the time-dependent data density at observation i , and $\varepsilon_i^2 = (\sigma_{ii}^2 + f_{ii}^2)/b_{ii}^2$ is the ratio of error variances for observation i , where σ_{ii}^2 is
 174 the observational error covariance, f_{ii}^2 is the interpolation error covariance, and b_{ii}^2 is the background error covariance interpolated to
 175 observation i . The background error covariance is estimated from the variability over time at each grid point of the original free-running
 176 simulation, and is not updated during the assimilation. For the radial velocity (and similarly for azimuthal velocity), Eq. (1) is expressed
 177 in incremental form for grid point k by (L91, Eq. 3.18)

$$\delta u_k = \lambda \sum_i \mu_{ki} \tilde{Q}_i R_i^2(\delta t_i) d_i \quad (3)$$

178 where the first two terms inside the sum come from Eq. (2), $R_i^2(\delta t_i)$ represents the weighting of each observation in time (Fig. 6),
 179 $d_i = y_i - H_i(\mathbf{u}^b)$ is the increment at observation i , and λ is a relaxation coefficient.

180 Observation i is weighted depending on the distance r_{ki} to grid point k using the expression

$$W_{ki} = \left(1 + \frac{r_{ki}}{s} \right) \exp \left(\frac{-r_{ki}}{s} \right) \quad (4)$$

Table 2. List of assimilations done. ΔT varies between 4.00 and 4.05 degC — see the SI for precise values.

Name	Datafile	Regime	Assimilated subset	Verifying subset	$\Omega/\text{rad s}^{-1}$	$t^{\text{start}}/\text{s}$	Length / s	Section
m2s1	expf2	2S	blue	red	0.775	8435	750	4.1
m2s2	expf2	2S	blue	red	0.800	6710	750	4.1
m3av2s1	expf2	3AV→2S	blue	red	0.825→0.800	5685	750	4.2
m3av1	expf2	3AV	blue	red	0.850	2135	750	4.2
m3av2	expf2	3AV	blue	red	0.875	435	750	4.3
m3sv2a	expf5	3SV	red	orange	2.3	2250	500	4.4
m3sv2b	expf5	3SV	blue, red	orange	2.3	2250	500	4.4
m3sv3a	expf5	3SV	red	orange	2.4→2.5	5350	500	4.5
m3sv3b	expf5	3SV	red	orange	2.4→2.5	5350	500	4.5
m3sv4	expf5	3SV	red	orange	2.4→2.5	5580	70	4.5
m3sv6	expf6	3SV	blue	red	2.8	4700	500	4.6
m3sv5	expf6	3SV	blue	red	3.0	9050	500	4.6
m3sv7	expf6	3SV	blue	red	3.1	10250	400	4.6

181 where s is a correlation scale, determined empirically to give the most accurate analysis when the analysis is compared with an
 182 independent observational dataset. We use both horizontal and vertical weighting functions to set

$$\mu_{ki} = W_{ki}^h W_{ki}^v \quad (5)$$

183 In the laboratory experiment, only horizontal velocity observations are available. Therefore balanced increments to the pressure and
 184 temperature fields are calculated using the geostrophic and hydrostatic approximations (except in the boundary layers, where these
 185 relations do not apply). The vertical velocity and the temperature and pressure in the boundary layers are updated only by the model.
 186 Each of the terms above is discussed in more detail in the SI.

187 4. Results

188 In this section we present a series of assimilations using ACACIA from the three observational datafiles, covering the whole range
 189 of regimes and rotation rates available. Each assimilation is between 400 and 750 s long, starting at $t^a = t^{\text{start}}$ and finishing at t^{stop} .
 190 The assimilation at t^{start} uses a 1850 s free-running MORALS run as its background state and observations between $t^{\text{start}} - t^f$ and
 191 $t^{\text{start}} + t^b$ as input to ACACIA, returning an analysis at t^a . t^f and t^b are the lengths of the assimilation window on either side of t^a
 192 (Fig. 6). MORALS then advances this analysis to $t^a + \Delta t$. The MORALS forecast at this new t^a becomes the background for
 193 the next ACACIA assimilation, using observations between $(t^{\text{start}} + \Delta t) - t^f$ and $(t^{\text{start}} + \Delta t) + t^b$ to compute the next analysis.
 194 And so on, until the final analysis at $t^a = t^{\text{stop}}$.

195 Each assimilation is done using one subset of the observations and verified against another. We measure the accuracy of each analysis
 196 using the RMS *residual error* between the analysis and the verifying observations. For N analysis-verification pairs at a particular time
 197 and vertical level, the residual error for radial velocity (similarly for v_{rmse}) is

$$u_{\text{rmse}} = \sqrt{\frac{1}{W} \sum_{n=1}^N w_n (u_n^o - H_n[\mathbf{u}^a])^2} \quad (6)$$

198 where $H_n[\mathbf{u}^a]$ is the analysis interpolated to the n th observation point, and u_n^o is the n th point in the verifying observations. The
 199 residual error is representative of all the analysis-observation pairs at a particular time, but the observations are irregularly distributed
 200 in space, so we assign each pair a weight w_n depending on the local data density (expression derived in the SI), with W the total weight
 201 over all observations.

202 If the assimilated and verifying datasets are independent (which they are assumed to be), then on average the minimum possible
 203 residual error between analysis and verifying observations is the observation error, 0.0057 cm s^{-1} . However, reaching this ‘perfect
 204 analysis’ would probably indicate data overfitting. At the other end of the scale, the residual error associated with a ‘climatological
 205 analysis’ just using the mean observed velocity everywhere would be the standard deviation of the observations away from their mean,
 206 and so sets an upper reference point for our residual error diagnostics. This standard deviation is 10–15 times the observational error,
 207 depending on the datafile and vertical level.

208 Table 2 summarises the different assimilations, and Fig. 2 shows the position of each assimilation within its datafile. **Young and Read**
 209 **(2008b)** found a systematic shift of -0.11 rad s^{-1} between observations and model in the position of the wavenumber 2 / wavenumber 3
 210 regime transition, and we applied this shift when setting the model’s rotation rate. The assimilations use an optimized set of parameters
 211 that were determined based on a series of short test assimilations. AC is usually optimized in an empirical sense, where the ‘best’
 212 parameter values are those which give the most accurate analysis. We measured the accuracy of the analysis in these tests using (a) the
 213 residual error (Eq. 6), (b) the number of assimilation cycles required for the observations and background to converge, indicated by
 214 the first minimum in the residual error curve, (c) the real time taken to compute each assimilation, and (d) a visual comparison of the
 215 analysis and observations, to note any systematic errors. The optimal parameters so obtained are listed in Table 3; details of the tests
 216 themselves are included in the SI.

Parameter	expf2	expf5 / expf6
$\Delta t / s$	2.5	1.0
t^f / s	26	21
t^b / s	26	21
s_{\min}^h / cm	0.21	0.1
s_{\max}^h / cm	0.42	0.2
s_{\min}^v / cm	0.5	0.75
s_{\max}^v / cm	0.75	1.00

Table 3. List of optimal parameters for the assimilations in the regular (expf2) and chaotic (expf5, expf6) flow regimes.

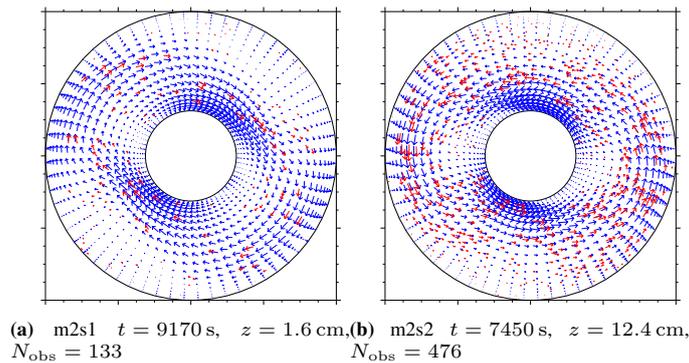


Figure 7. Horizontal velocity field analysis (blue, grey in print) and verifying observations (red, black in print) for two cases at different levels in the 2S assimilations.

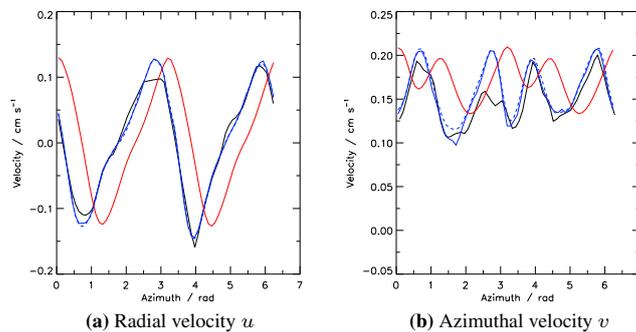


Figure 8. Azimuthal sections at $R = 5.25 \text{ cm} / z = 9.7 \text{ cm}$ at $t = 9155 \text{ s}$ (final assimilation) in assimilation m2s1. The four lines are the analysis (solid), background (dashed), both of which are blue in the online version and dark grey in print, observations interpolated to the same azimuthal section (black), and free-running run (red, light grey in print).

217 We also ran a free-running MORALS simulation parallel to the sequence of assimilations, starting from the transformed background
 218 state at $t = t^{\text{start}}$ and free-running afterwards, using the same Ω and ΔT forcing parameters as the assimilation itself. The only
 219 difference between this and the analysis was that the assimilation step was omitted. Including this free-running simulation for
 220 comparison, we can see how fast the model and analysis would otherwise diverge, the difference the assimilation step makes to the
 221 model state, and any systematic differences between model and analysis.

222 4.1. 2S assimilations: m2s1 and m2s2

223 The first two assimilations tracked the observations over two 750 s segments in the 2S flow regime, where the main aim is to track the
 224 position of the baroclinic wave.

225 Figures 7 and 8 show typical analyses, and are good fits to the verifying observations. The assimilation converges to an optimum
 226 analysis (i.e. minimum residual error) in less than 25 s. Throughout the assimilations they compare well with the verifying observations,
 227 with residual error shown in Fig. 9. In both assimilations the accuracy of the analyses decreases from the top to the bottom of the tank,
 228 presumably because there are more observations at the top.

229 In the free-running simulations, the model runs are most accurate at $z = 7.0 \text{ cm}$. This is probably because the background
 230 transformation matches up the model and observations at 7.0 cm before starting the run. The residual error slowly increases over
 231 time instead of remaining approximately constant (like the analysis residual errors) because the modelled and observed waves drift out
 232 of phase over time but are not pulled back together by the assimilation.

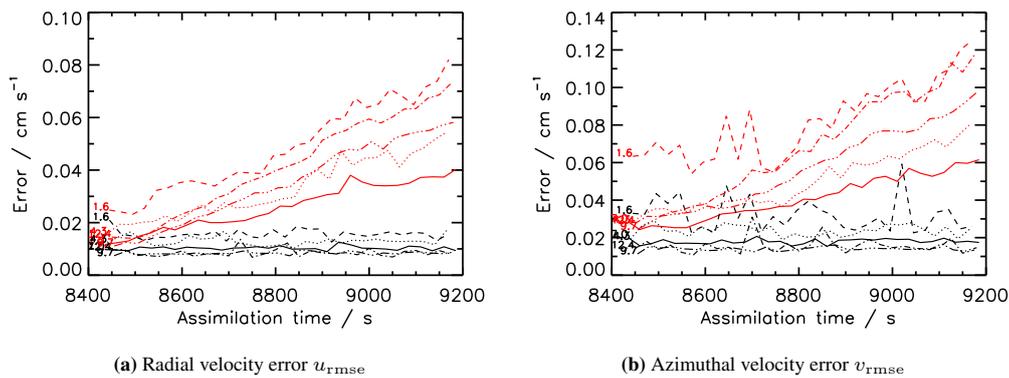


Figure 9. Analysis (black) and free-running model (red, grey in print) residual error (Eq. 6) at each of the five observation levels over assimilation m2s1. The residual error for run m2s2 was very similar. The five lines correspond to $z = 12.4$ cm (dot-dash), 9.7 cm (dot-dot-dot-dash), 7.0 cm (solid), 4.3 cm (dotted), and 1.6 cm (dashed).

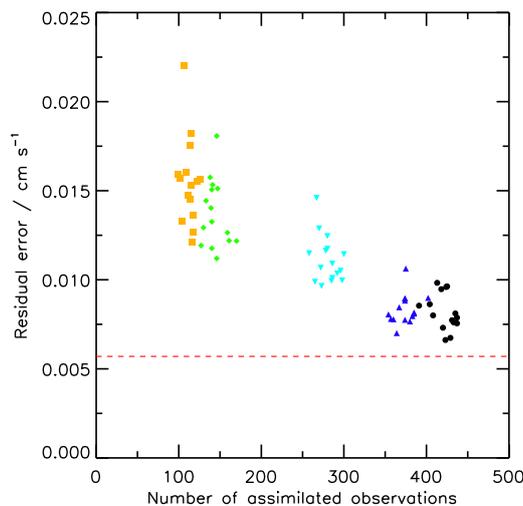


Figure 10. Analysis residual error for u in assimilation m2s2 against the number of observations used in each assimilation. The circles are from level 12.4 cm, upwards triangles are 9.7 cm, downwards triangles are 7.0 cm, diamonds are 4.3 cm, and squares are 1.6 cm. The horizontal dashed line indicates the observational error 0.0057 cm s^{-1} .

233 The analysis residual error in Fig. 9 never falls to the theoretical minimum set by the observational error, although it gets close. The
 234 error does decrease as the number of observations increases, however. This suggests there may be an observation number for which
 235 the residual error matches the observational error. This number can be estimated by plotting the residual error for each analysis over
 236 a sequence of assimilations against the number of observations in each assimilated dataset. This is shown in Fig. 10 for assimilation
 237 m2s2, showing that about 500–600 observations would be required for u_{rms} to fall to the minimum.

238 4.2. 3AV assimilations: m3av1 and m3av2

239 In the 3AV regime the assimilation aims to track both the position of the baroclinic wave and its vacillation over time. Figure 11 shows
 240 typical results, and again the analyses show a close correspondence between analysis and verifying observations, with residual errors
 241 similar in magnitude to the 2S case. There is no oscillation in the analysis residual error over the vacillation timescale, but in the
 242 free-running runs there is a clear oscillation, presumably because the vacillation period in the model and observations do not match
 243 exactly.

244 4.3. 3AV to 2S transition assimilation: m3av2s1

245 The final case using datafile expf2 was to assimilate over a regime transition from 3AV to 2S as Ω is reduced from 0.825 to
 246 0.800 rad s^{-1} . The aim is to assimilate first the abrupt change in Ω at $t = 5750$ s, followed by the wavenumber change between
 247 $t = 6000$ s and 6300 s, and to track the position and vacillation of the original wave as well as its change from a vacillating to non-
 248 vacillating wave.

249 Figure 12 shows a sequence of analyses and azimuthal slices at $z = 9.7$ cm between $t = 5905$ s and $t = 6405$ s. ACACIA performs
 250 as well in this scenario as it does in the previous cases. This figure also shows the wavenumber transition in the free-running model.
 251 The transition happens about 100 s before it happens in the analyses, and once it has occurred in the free-running model it is almost
 252 perfectly in phase with the analysis, but this is coincidental.

253 The v residual error is shown in Fig. 13. Like the other assimilations, the analysis residual remains approximately constant
 254 throughout the assimilation. There is no rise in the residual error as the wavenumber changes during the later part of the assimilation,
 255 but there is a small jump at the time the rotation rate changes at 5750 s, when observations with the new Ω first enter the assimilation

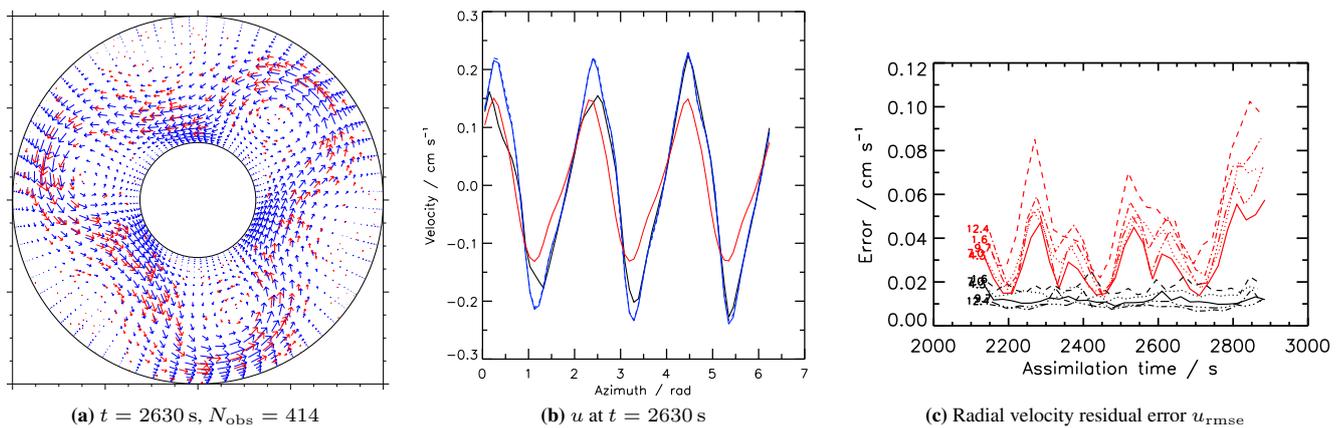


Figure 11. 3AV results from assimilation m3av1. Left: Horizontal velocity analyses and verifying observations at $z = 9.7$ cm (colour key as Fig. 7). Middle: Azimuthal velocity sections through $R = 5.25$ cm, $z = 9.7$ cm at $t = 2630$ s (as Fig. 8). Right: Radial velocity residual error u_{rmse} (as Fig. 9). Assimilation m3av2 is qualitatively similar.

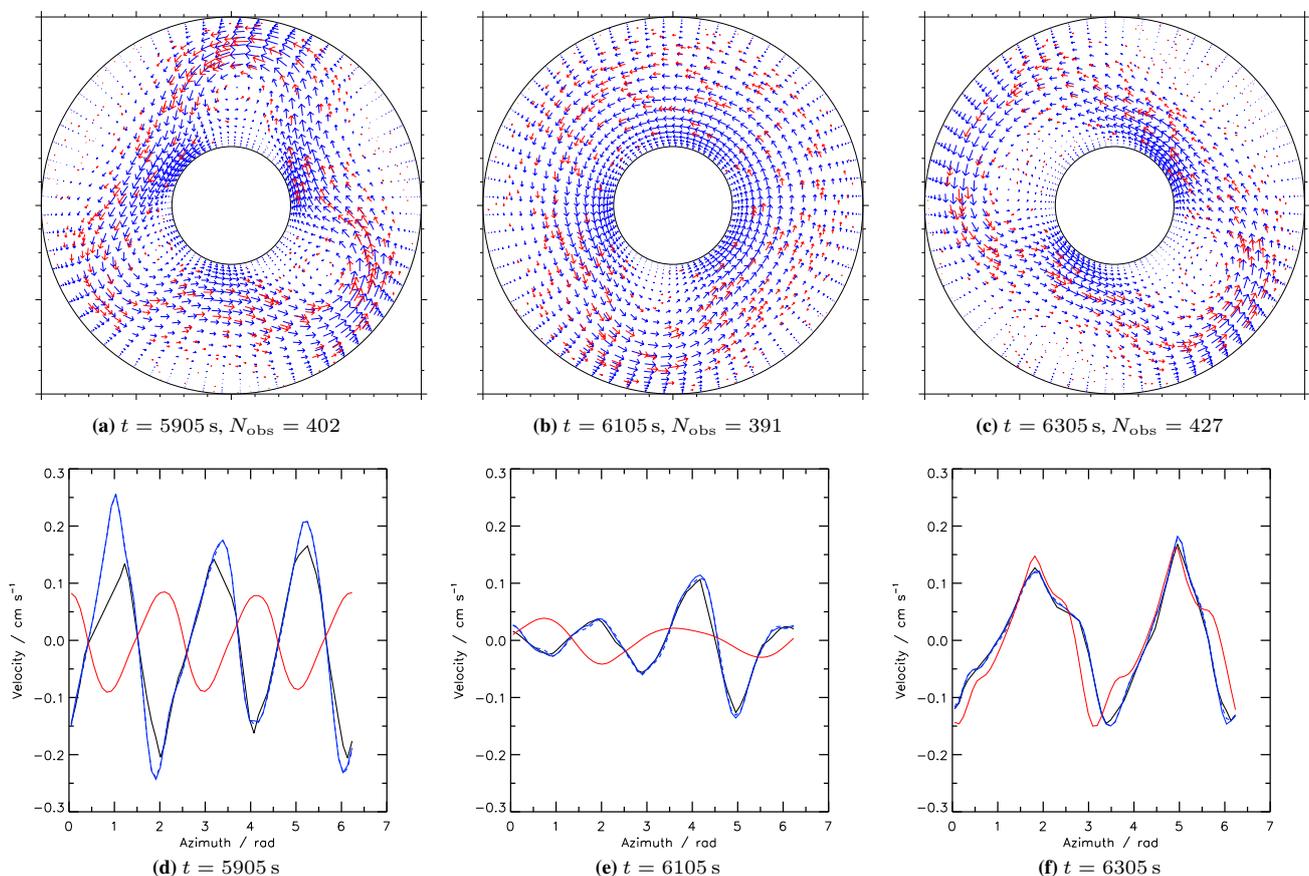


Figure 12. Top: Horizontal velocity analyses and verifying observations for three analyses spanning the wavenumber transition in assimilation m3av2s1, at $z = 9.7$ cm, frames separated by 200 s (colour key as Fig. 7). Bottom: corresponding radial velocity azimuthal sections at $R = 5.25$ cm (as Fig. 8). The SI contains these two sequences as animations.

256 window. This only occurs in v because the immediate effect of changing Ω is a sharp change in the azimuthal velocity in the rotating
 257 frame of reference, while the radial velocity is continuous in the rotating frame across the drop in Ω . There will be some boundary
 258 layer effects on the radial velocity from changing Ω , but these are minimal.

259 4.4. 3SV assimilations at low rotation rate: m3sv2a, m3sv2b

260 In the 3SV regime the aim is to capture the shape vacillation of the baroclinic wave as well as any chaotic dynamics. This second
 261 condition makes assimilation in the 3SV regime somewhat more difficult than in the 2S and 3AV regimes.

262 Figure 14 shows a representative horizontal velocity analysis at the upper level along with azimuthal sections at mid-radius from
 263 assimilation m3sv2a ($\Omega = 2.3$ rad s $^{-1}$). As in the previous sections, the analysis generally tracks the verifying observations well.

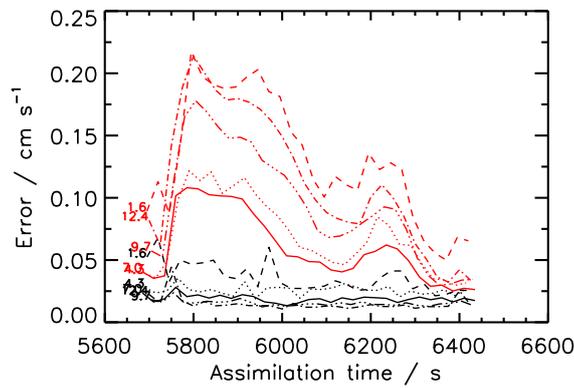


Figure 13. As Fig. 9, but for azimuthal velocity residual error v_{rmse} in assimilation m3av2s1.

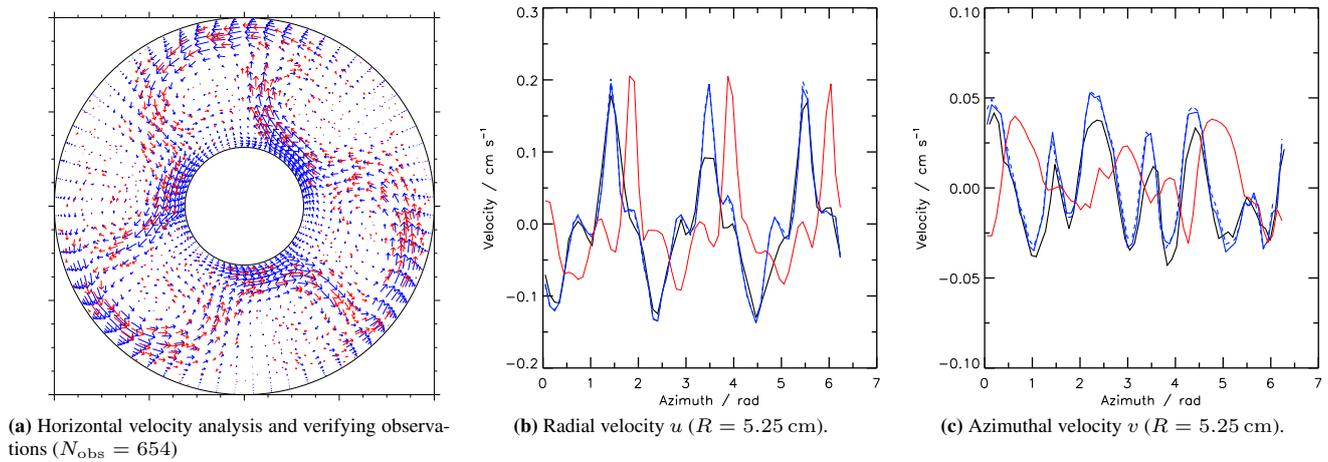


Figure 14. Representative assimilation m3sv2a 3SV horizontal velocity analysis and velocity sections at $t = 2750$ s, $z = 9.7$ cm. The line colours and styles are as in Figs 7 and 8.

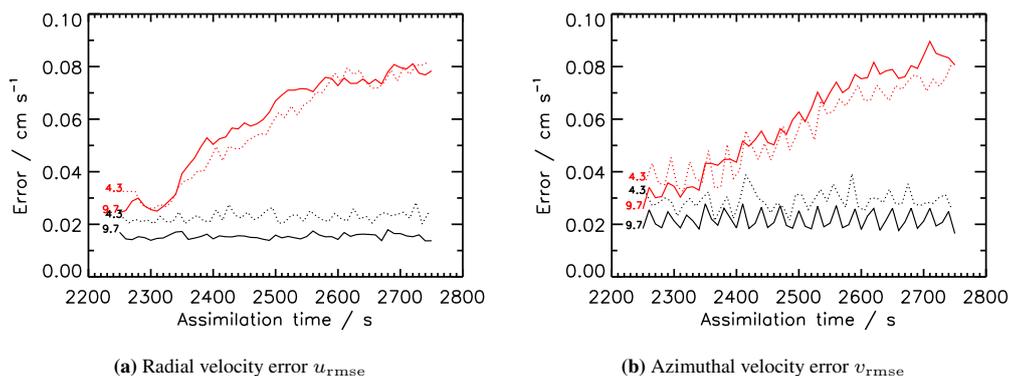


Figure 15. Analysis and free-running simulation residual errors at the two observation levels during assimilation m3sv2a. The solid lines correspond to $z = 9.7$ cm, and the dotted lines to 4.3 cm. Colours are as Fig. 9.

264 Figure 15 shows the residual error for m3sv2a. The error in the free-running simulation (red lines) slowly increases over time,
 265 primarily because the background state drifts out of phase with the observations (see the azimuthal sections in Fig. 14). The residual
 266 error is about 50% larger than for regular flow, and there is an oscillation in v_{rmse} with a steady period of 30 s, for which no explanation
 267 has been found (and was constant for different assimilation parameters).

268 Figure 10 suggested that the analysis residual error depends on the number of observations in the assimilated dataset, and that with
 269 a sufficiently high number of assimilated observations we might get a residual error comparable with observational error. To test this
 270 we ran assimilation m3sv2b, identical to m3sv2a but using both blue and red subsets for assimilation instead of just red (verification
 271 used subset orange, as before). Figure 16 shows the change in the residual error between m3sv2a and m3sv2b. Doubling the number
 272 of assimilated observations gives only a marginal improvement. This suggests that subsets blue and red are not as independent as
 273 first thought. Some observations remain within the light sheet between subsets and so appear in more than one, so the number of
 274 independent assimilated observations in m3sv2b is not double that in m3sv2a. Hence it is not possible to improve the accuracy much

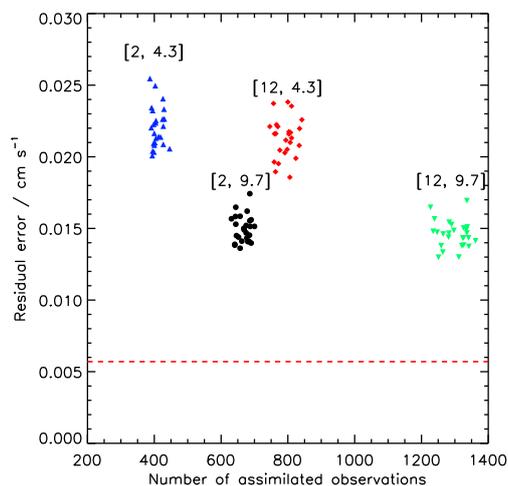


Figure 16. Radial velocity analysis residual error against the number of observations used in each assimilation, for assimilations m3sv2a and m3sv2b. The assimilated subset(s) (1-blue and 2-red) and vertical level are indicated in brackets, and the dashed line is the observational error 0.0057 cm s^{-1} .

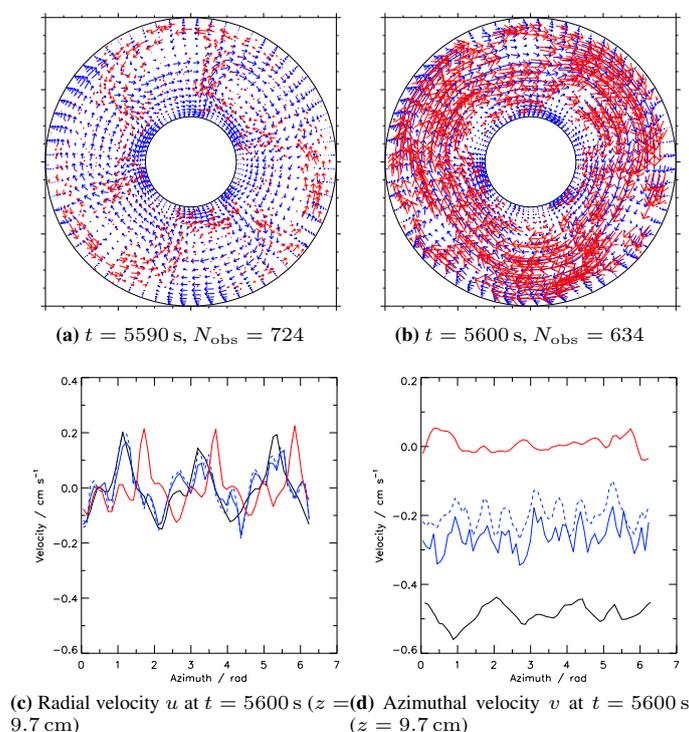


Figure 17. Top: Horizontal velocity analyses and verifying observations in assimilation m3sv3a, for two consecutive observation datasets at $z = 9.7 \text{ cm}$ (colours as Fig. 7). Bottom: azimuthal sections at $R = 5.25 \text{ cm} / z = 9.7 \text{ cm}$ at $t = 5600 \text{ s}$ (colours as Fig. 8).

275 by increasing the number of assimilated observations in this way. One could randomly split a single subset into two parts and use half
276 for assimilation and half for verification, but we leave this for the future.

277 4.5. 3SV assimilations over a rotation rate transition: m3sv3a, m3sv3b, m3sv4

278 Using the next set of assimilations we investigated how well ACACIA could assimilate data over an abrupt Ω transition in the
279 observational data. These assimilations began at $\Omega = 2.4 \text{ rad s}^{-1}$ and stepped up to $\Omega = 2.5 \text{ rad s}^{-1}$ at $t = 5600 \text{ s}$ (Fig. 2). The
280 observations show a short burst of high azimuthal velocity at the transition before the wave settles back down. After the transition
281 the wave settles down to almost exactly the same place in the tank, showing that the particles gathered on the bottom of the tank have
282 a strong feedback effect on the position of the wave, and do not disperse when the flow temporarily speeds up.

283 Figure 17 shows two analyses just before and at the Ω transition in m3sv3a, along with azimuthal sections at the transition. We see
284 the effect of the transition on the analysis 10 s before it takes place in the observations (Fig. 17a), as the assimilation window begins to
285 include observations from beyond the transition. When Ω changes (Fig. 17b) the observed flow changes over a timescale much shorter
286 than both the time between observations and the assimilation window. As a result the increased v enters the window before the flow

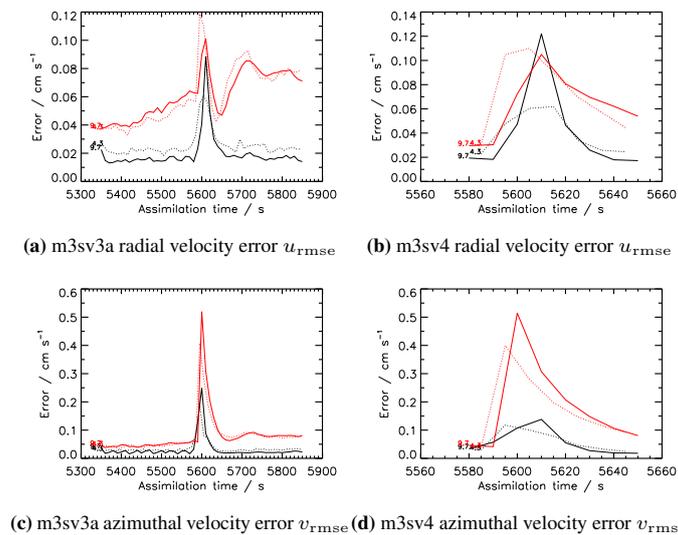


Figure 18. Left: Analysis and free-running simulation residual errors at the two observation levels during assimilation m3sv3a. Colours and line styles are as Fig. 15. Right: As left column but for assimilation m3sv4.

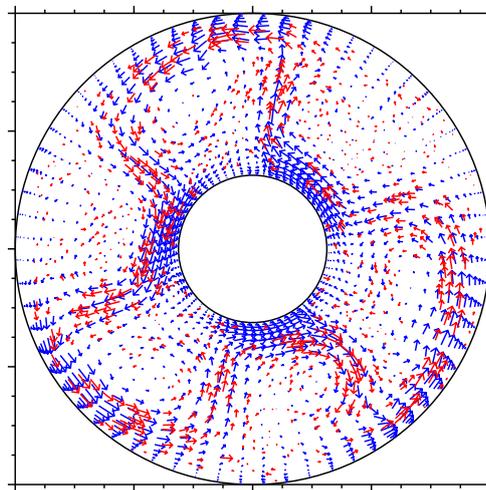


Figure 19. Representative high- Ω analysis, from m3sv5 at $t = 9550$ s and $z = 9.7$ cm, showing the horizontal velocity field (blue, grey in print) and verifying observations (red, black in print), with $N_{\text{obs}} = 731$.

287 is forced by the model in a way that produces such a flow dynamically, so the quality of the assimilation is severely degraded. This is
 288 reflected in the residual error (Fig. 18a,c), with the errors increasing by a factor of six in u_{rmse} and ten in v_{rmse} over 10 s.

289 To try to improve ACACIA's performance we re-ran the transition scenario with a reduced assimilation window using $t^{\text{b}} = t^{\text{f}} = 11$ s
 290 (m3sv3b), and then also a reduced time of $\Delta t = 0.2$ s between assimilations (m3sv4). In m3sv3b observations from beyond the
 291 transition would enter the window at a time closer to the transition itself, but the improvement at the transition was marginal. With the
 292 reduced time between assimilations in m3sv4 there was still an increased residual error during the transition (Fig. 18b,d), but there was
 293 a 40% improvement in v_{rmse} over m3sv3a (no improvement in u_{rmse}). This improvement is due to there being more iterations using
 294 observations *after* the transition *before* the transition actually occurs.

295 4.6. 3SV assimilations at high rotation rate: m3sv5, m3sv6, m3sv7

296 These final three assimilations were at the higher end of the available rotation rates, up to 3.1 rad s^{-1} in m3sv7*. As Ω increases over
 297 dataset expf6, vortices begin to be shed westward from the cyclones at the upper level, becoming more and more pronounced as Ω
 298 increases; the aim here is to assimilate this behaviour accurately.

299 Figure 19 shows a representative analysis from one of these assimilations and Fig. 20 the residual error. The analyses are generally
 300 accurate, even at the highest Ω . The v residual error oscillates substantially, unlike the lower- Ω results, but this is the only substantive
 301 difference. This may be because the model and observations have slightly different wave drift periods, which will introduce errors into
 302 the v (less so the u) velocity analyses. The oscillation period is 150–200 s, however, which does not match any known timescale in this
 303 flow regime.

*For these assimilations we returned the parameters to the defaults listed in Table 3.

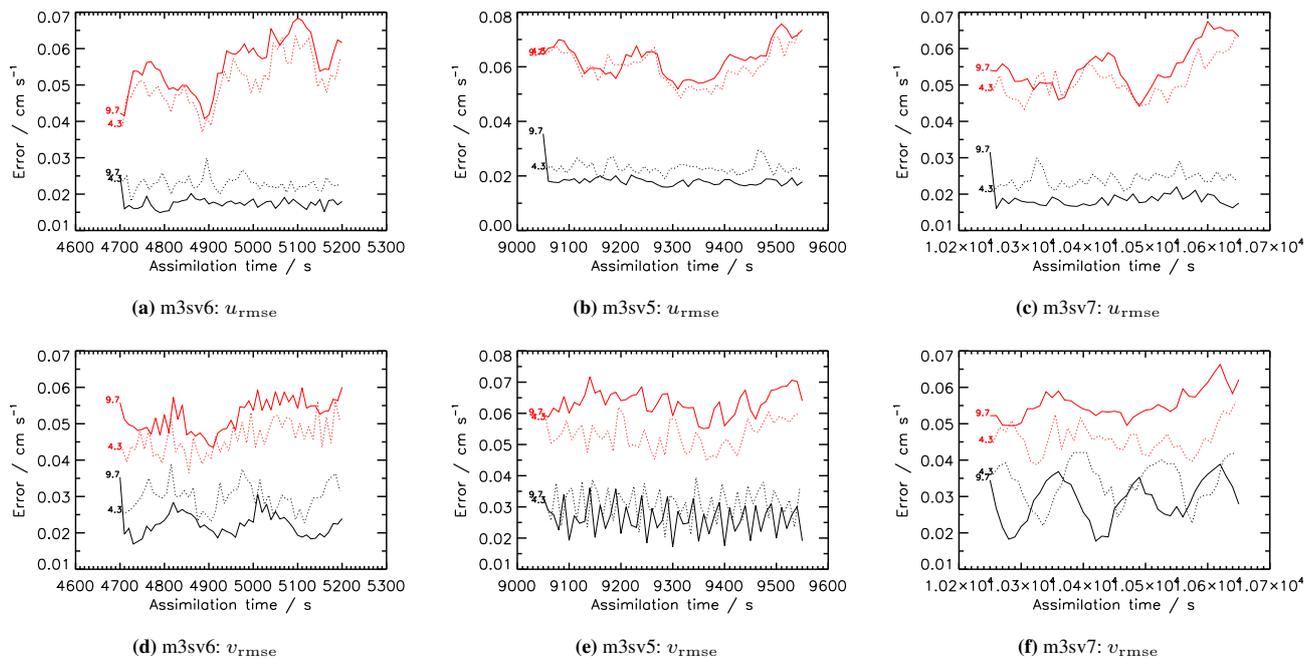


Figure 20. Radial and azimuthal residual errors in assimilations m3sv6, m3sv5, and m3sv7 (increasing Ω). Line colours and styles are as Fig. 15.

304 Vortex shedding is reasonably well assimilated by ACACIA, and is shown at the upper level in Fig. 21 for a case in m3sv7. The
 305 vortices move from the westward edge of the cyclone into the anticyclonic region over a timescale of 100–150 s. In the next section we
 306 examine these vortices in more detail.

307 4.7. Comparisons between chaotic and regular assimilations

308 Figure 22 combines the results from all the assimilations to show the distribution of residual errors at the different vertical levels. It
 309 shows clearly that u error is generally smaller than v error, that the error decreases as the vertical position increases, as expected from
 310 the dependence on the number of observations with height, and that the assimilations were more accurate in the regular regimes than in
 311 the chaotic regime. The best analyses came close to the observational error. No analyses approached the variability in the observations
 312 above 0.06 cm s^{-1} representative of a ‘climatological analysis’, which is encouraging.

313 Figure 23 compares residual errors between regular and chaotic flow as well as within regimes. Clearly the flow regime has the
 314 largest effect on the values: assimilations in the chaotic regime (3SV) have a consistently higher error than those in regular regimes (2S
 315 and 3AV). Within a particular regime there is little variation among the different assimilations.

316 5. Probing annulus dynamics with the ACACIA analyses

317 In the chaotic regime we see vortex shedding from the cyclones at high Ω (Fig. 21), but we don’t see this at low Ω . Is there a mechanism
 318 that causes vortex shedding at high rotation but not at low rotation? We are only able to assimilate horizontal velocity observations, but
 319 the assimilation retrieves unobserved quantities such as temperature, which we can use to study this phenomenon more closely.

320 Figure 24 shows representative temperature fields at $z = 9.7 \text{ cm}$ in four cases: regular assimilations m2s1 and m3av2 at 0.775 and
 321 0.875 rad s^{-1} , and chaotic assimilations m3sv2a and m3sv7 at 2.3 and 3.1 rad s^{-1} . Temperature is retrieved by the assimilation using
 322 the hydrostatic and geostrophic approximations. There is a clear difference between the temperature structure in the regular and chaotic
 323 flow regimes. Regular flow contains fronts between low and high temperature regions, but in the chaotic flow regime the cool fluid is
 324 constrained to a radial tongue or ‘plume’ along the westward side of each cyclone (using the streamfunction to identify the cyclone’s
 325 position). The streamfunction shows why this occurs. In the chaotic flow near the inner cylinder within the cyclones the main jet
 326 is distorted so that the streamlines almost fully isolate the cyclone from the inner cylinder. Hence cool fluid can only be channelled
 327 outwards via the plume, rather than diffusing into the interior of the cyclone. By contrast, in the regular flow the main jet is not distorted
 328 near the inner cylinder. Hence cool fluid more easily diffuses into the body of the cyclone. The strong jet then isolates the cool region
 329 from the warm region, maintaining the front.

330 We can demonstrate what happens to cool fluid as it is advected outwards by the jet in the chaotic flow using the potential energy
 331 (E_P) field, where $E_P = \Delta\rho g z$ with $\Delta\rho$ the departure from a reference density. High temperature corresponds to low E_P and vice
 332 versa. In the chaotic flow there is a region of increased E_P extending into the anticyclonic region (c.f. Fig. 24d), which is more
 333 pronounced at the high rotation rates. In the chaotic flow pulses of cool fluid are periodically advected by the jet towards the outer
 334 cylinder, the pulse being stronger at higher Ω . This modulates the cool tongue described above over a $\sim 100 \text{ s}$ period. We include an
 335 animation of the m3sv6 E_P with the SI showing this phenomenon. When the pulse reaches the outer cylinder it is either carried around
 336 the cyclone by the flow (e.g. $t \approx 4920 \text{ s}$ in the animation) or is smeared out in the azimuthal direction (e.g. $t \approx 5200 \text{ s}$); little E_P seems
 337 to be carried to the outer cylinder itself.

338 Is the jet being forced in different ways by the eddies at the different rotation rates, and can we use this to explain why vortex
 339 shedding occurs in one case and not the other? We consider two possibilities here that may help us to understand the phenomenon.

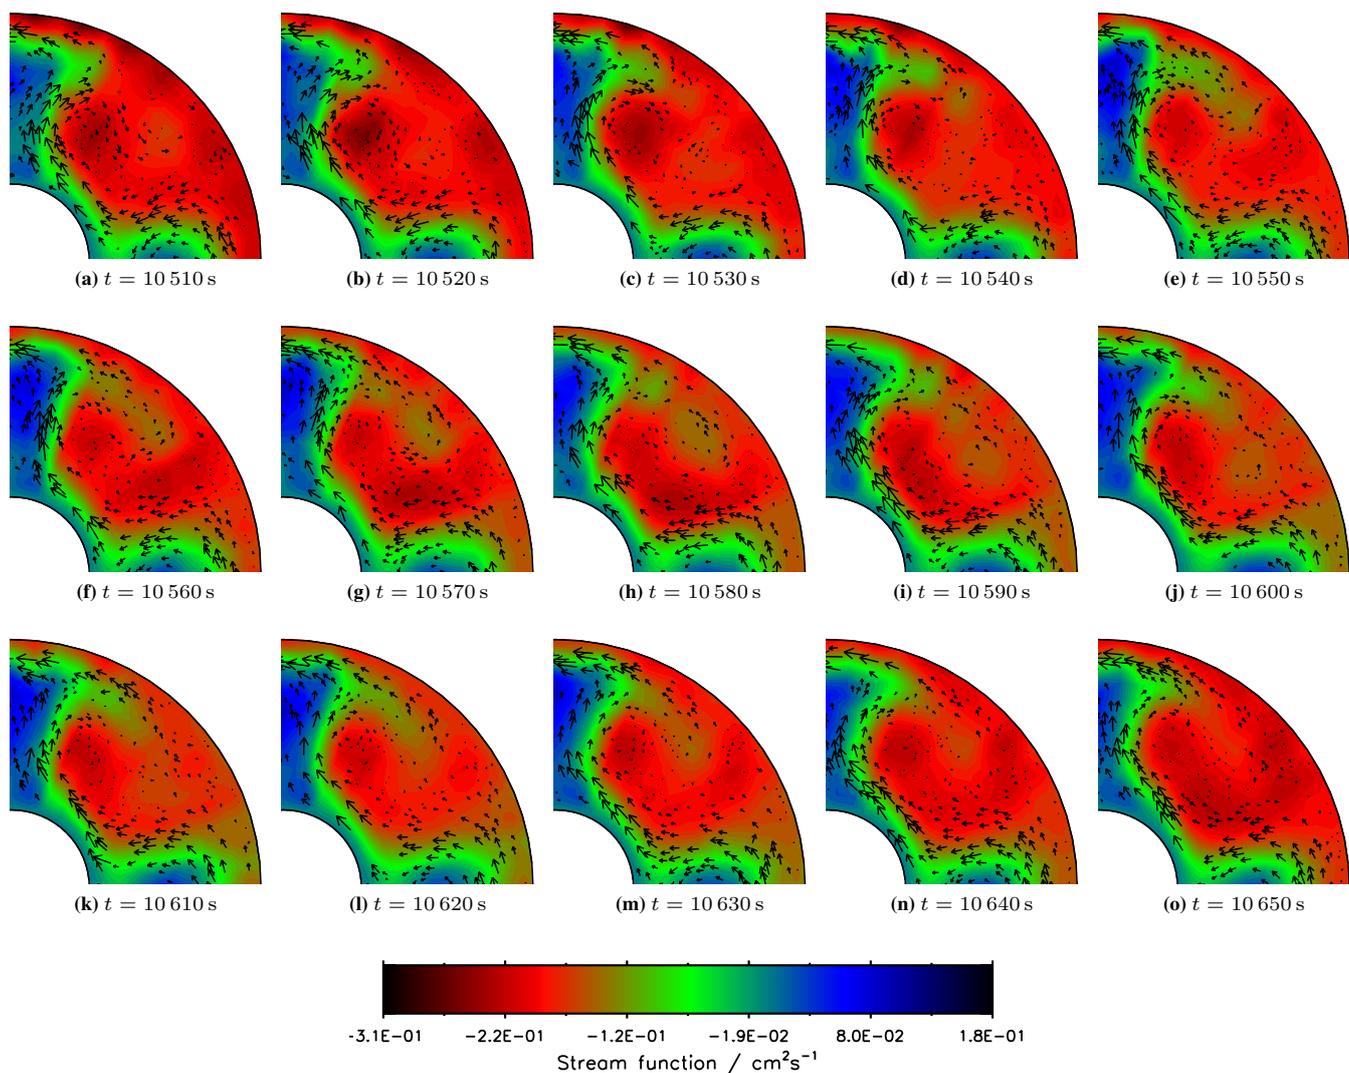


Figure 21. Sequence of analyses at $z = 9.7$ cm over 140 s in assimilation m3sv7, showing the development of a vortex and its shedding from the westward edge of a cyclone in the upper right quadrant of the annulus. The analysis is shown as horizontal stream function contours, with the observed velocity vectors superimposed (colour online only).

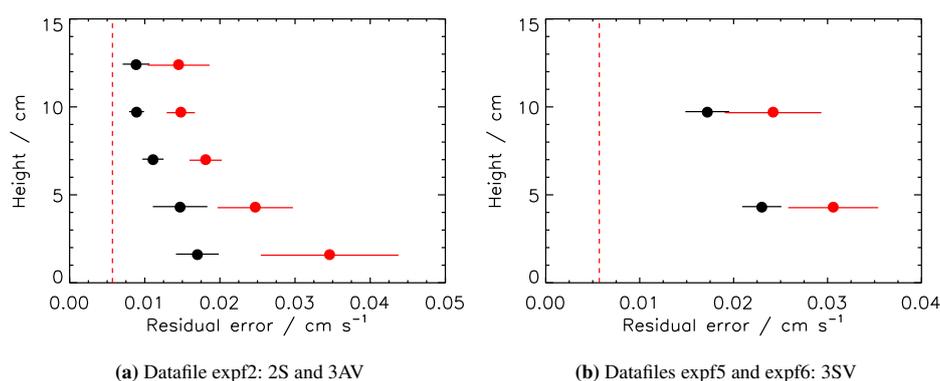


Figure 22. Distribution of u (black) and v (red, grey in print) residual error for all analyses at each vertical level. The mean (dot) and one standard deviation (line) of the residual error over all the analyses at each vertical level are plotted. The vertical dashed line represents the observational error 0.0057 cm s^{-1} .

340 5.1. E-vectors

341 E-vectors identify regions where barotropic eddies stimulate both cyclonic and anticyclonic cyclogenesis. James (1994, Eq. 7.20b)
342 defines the E-vectors for atmospheric flow; redefining them in the coordinate system used by MORALS gives

$$\mathbf{E} = (E_R, E_\phi) = (-\overline{u'v'}, \overline{u'^2 - v'^2}) \quad (7)$$

343 where u and v are the assimilated radial and azimuthal velocities, primes denote eddy quantities with respect to the time mean over
344 the whole assimilation sequence, and the bar denotes a barotropic (vertically-averaged) time mean over the whole sequence. A whole

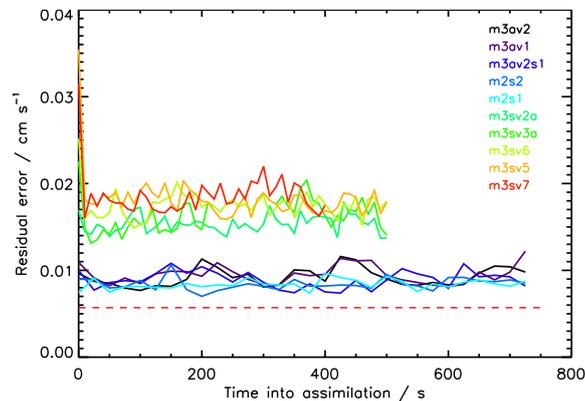


Figure 23. Comparison of radial velocity residual errors at $z = 9.7$ cm for all assimilations. In the online version, red, green, and yellow lines are from 3SV assimilations, and blue and black lines are from 2S and 3AV. The colour key is shown on the plot, along with the observational error as a dashed line. The Ω -transition part of m3sv3a has been omitted.

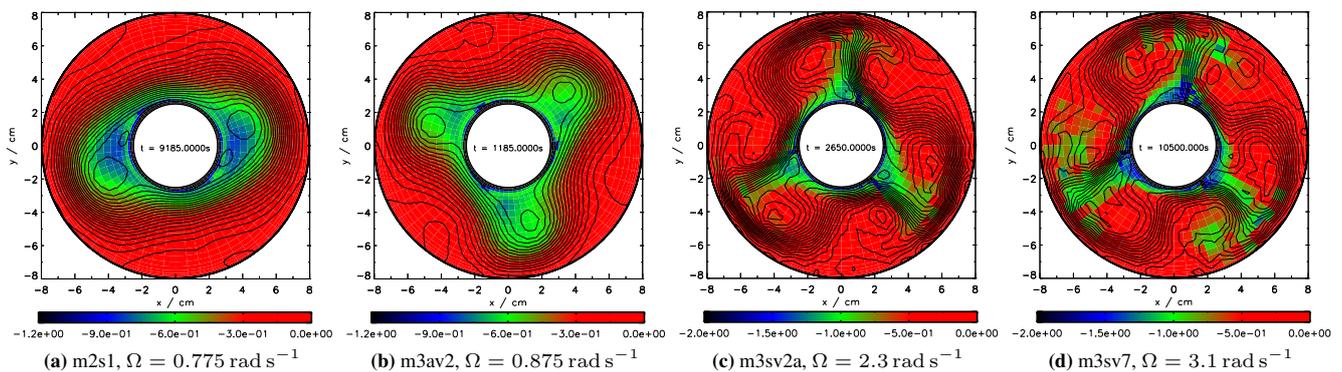


Figure 24. Representative assimilated temperature fields at $z = 9.7$ cm for regular (left two) and chaotic flow (right two). Temperatures are relative to 22°C . Assimilated horizontal streamfunction contours are plotted on top. Note the colour scale is slightly different for the chaotic and regular cases (colour online only).

345 assimilation sequence is hence represented by a single horizontal section. We can calculate the barotropic E-vectors only because
 346 the assimilation fills in the horizontal velocity field at unobserved heights. James (1994, Fig. 7.15) shows that for E-vectors pointing
 347 towards the east (i.e. in the positive azimuthal direction in the annulus), there is anticyclonic cyclogenesis southward of the divergent
 348 region (i.e. nearer the outer cylinder), cyclonic cyclogenesis southward of the convergent region, and vice versa northward of these
 349 regions.

350 Figure 25 shows the E-vectors and their divergence $\nabla \cdot \mathbf{E}$ for three chaotic assimilations: m3sv2a, m3sv6, and m3sv7. There is a
 351 clear trend in the strength of the divergence and convergence of the E-vectors as rotation increases. The divergent region within the
 352 anticyclone becomes larger and stronger, as does the convergent region within the cyclone. This acts to reinforce the main cyclone
 353 nearer the outer cylinder, but weaken the part of the cyclone that extends into the anticyclonic region near the outer cylinder. (This
 354 extension is not clear in the barotropic time-averaged streamfunction, but see Fig. 24c for example; at $z = 9.7$ cm the extended region
 355 becomes more pronounced as rotation increases). This weakening of the cyclonic region extending into the anticyclonic region may
 356 then cause vortices to break off into the anticyclonic region.

357 5.2. Rossby wave stability

358 The onset of small-scale cyclic or chaotic time-dependence in the large-scale wave could represent a secondary instability, akin to
 359 the instability of the ‘classical’ Rossby wave studied by Hoskins (1973) and Gill (1974). This secondary instability can take either a
 360 barotropic or baroclinic form, depending upon conditions (see Grotjahn (1984) or Read (1993) for reviews).

361 Deininger (1982) showed that the nonlinear development of barotropic instability of a (barotropic) Rossby wave could lead to a
 362 periodic exchange of energy between the main wave and a perturbation, producing a phenomenology akin to ‘structural vacillation’.
 363 However, our flows do not appear to violate the barotropic stability criterion set out by Bell (1989) (for the quasi-geostrophic
 364 case, at least). He found $r^2(K^2 - 1) \leq 3$ to be sufficient for wave stability, where K is the wavenumber and $r = 2L/L_x$ is the
 365 horizontal aspect ratio with L the lateral length scale and L_x the zonal length scale. For an annulus of cylinder radii a and b we have
 366 $r = (2b - 2a)/(\pi b + \pi a)$, so in our case $r^2(K^2 - 1) \approx 0.9$, which is sufficient for stability.

367 Kim (1978) suggested that a baroclinic Rossby wave may be unstable to baroclinic secondary instabilities if the large scale
 368 wavelength L is much larger than the deformation radius $L_D = NH/f$, where N is the buoyancy frequency, H is a vertical length
 369 scale, and f is the Coriolis parameter. It is reasonable to speculate that nonlinear development of such an instability could lead to cyclic
 370 vacillations, much as for barotropic instabilities.

371 Our simulations suggest that Kim’s criterion for baroclinic secondary instabilities is satisfied. The horizontal wavelength for
 372 wavenumber-3 flow in our annulus is about 11 cm. For assimilation m3sv7 a typical buoyancy frequency $N = ((g/\rho_0)\partial\rho/\partial z)^{1/2}$
 373 is about 0.2 s^{-1} , which gives a deformation radius of $L_D \approx 0.5$ cm using $f = 2\Omega$ and the annulus depth as a vertical length scale.
 374 Note that calculating N requires assimilation in order to retrieve $\partial\rho/\partial z$, which is a function of temperature. Hence $L \gg L_D$, with

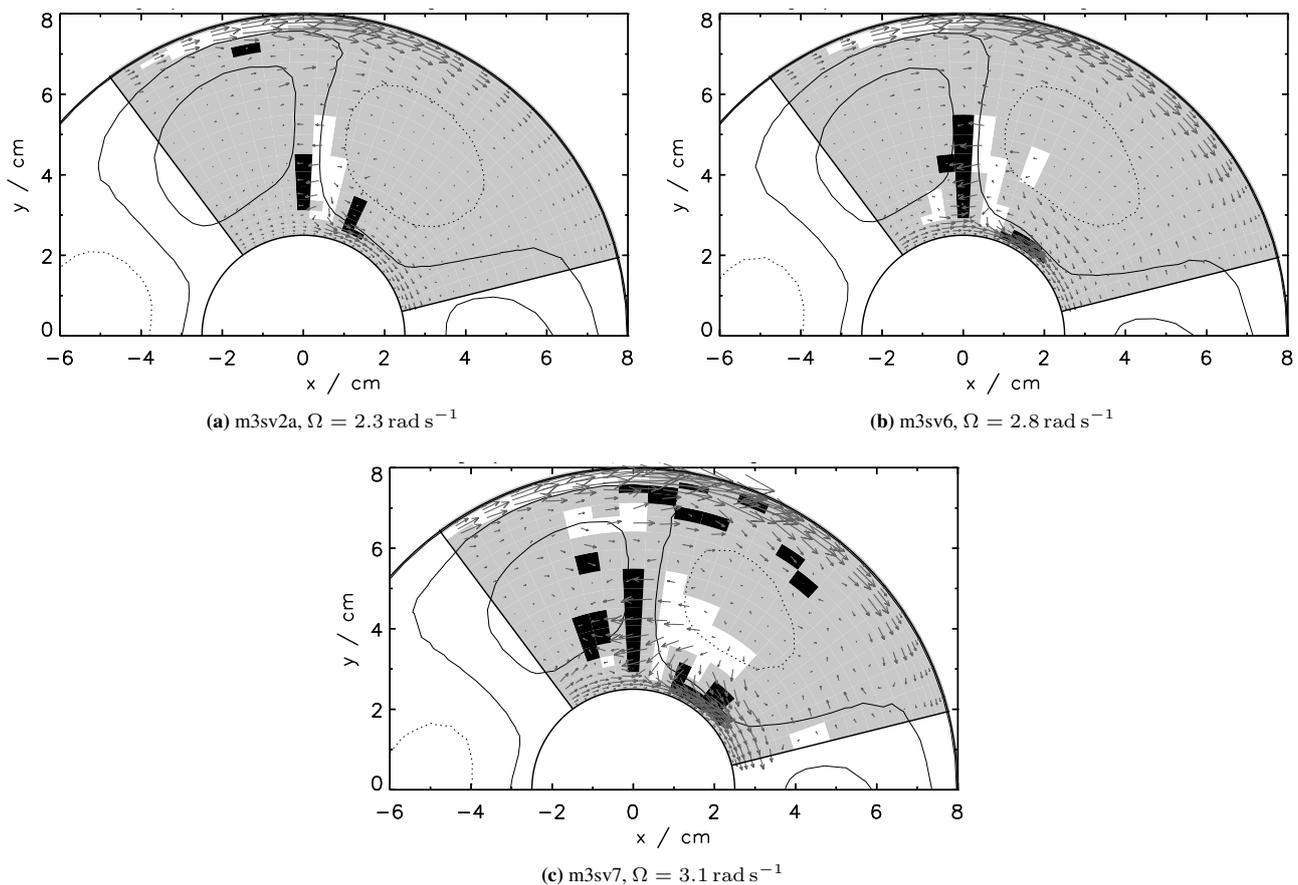


Figure 25. E-vectors for chaotic flow at low, medium, and high rotation rates. Each plot shows a segment of the annulus; the pattern is qualitatively similar for the other two waves. The black contour lines show the assimilated barotropic time-averaged horizontal streamfunction (contours below the middle of the range are dotted), and the grey vectors are the barotropic time-averaged E-vectors. The shading shows the E-vector divergence: black is up to $-5 \times 10^{-4} \text{ cm s}^{-2}$, grey is between $-5 \times 10^{-4} \text{ cm s}^{-2}$ and $+5 \times 10^{-4} \text{ cm s}^{-2}$, and white is above $+5 \times 10^{-4} \text{ cm s}^{-2}$. The colours, vectors, and contours use the same scale for each plot.

375 increasing supercriticality as rotation rate increases because $L_D \propto 1/\Omega$. Intriguingly, the horizontal scale of the vortices generated at
 376 high rotation rates (e.g. in Fig. 21) is similar to L_D .

377 This indicates a plausible interpretation for the eddy shedding observed in our assimilations. The fact that the small-scale eddies
 378 occur close to the outer cylinder, where thermal gradients and fronts are strongly evident, is consistent with this interpretation.

379 6. Conclusions

380 In this paper we have developed a method for data assimilation in the rotating annulus experiment using analysis correction (Lorenc
 381 *et al.* 1991). We used the algorithm, ACACIA, to assimilate laboratory data into a rotating annulus model to produce sequences of
 382 analyses in steady and amplitude vacillating regular flow regimes, and in a structurally vacillating chaotic flow regime. While this was
 383 not the first study to assimilate data successfully into a rotating annulus simulation, it was the first to examine specific flow regimes,
 384 and the first to use irregularly-distributed observations — a situation more realistic for operational meteorology.

385 Using an optimal set of parameters determined by a series of test assimilations, we demonstrated that ACACIA combines model and
 386 experimental data accurately in the rotating annulus context. In regular flow regimes, residual errors were between 1.5 and three times
 387 the observational error for radial velocity, and between three and seven times for azimuthal velocity, with error generally decreasing
 388 towards the top of the tank (Fig. 22). In chaotic flow the residual errors were generally higher, but still only 3–7 times the observational
 389 error, depending on the vertical level and the particular velocity component. All residual errors were significantly better than the 10–15
 390 times observational error that would result from an analysis containing only the mean observed velocities. There were no statistically
 391 significant correlations between residual error and Ω .

392 Assimilation across a wavenumber transition between 3AV and 2S proved to be as accurate as assimilation without such a transition.
 393 In the chaotic regime the algorithm also accurately propagated information from data-rich to data-poor areas. Assimilating over a
 394 rotation rate transition was not as accurate, but we were able to reduce the error significantly by tuning the assimilation parameters.

395 We also used the results from the assimilations to study the dynamics of the annulus experiment in a way that would not be possible
 396 without the assimilation. The temperature field was used to show that pulses of cool fluid are advected by the eastward jet into the region
 397 close to the outer cylinder, where they are smeared out in the azimuthal direction. The assimilation fills in the horizontal velocity field at
 398 unobserved heights and this allowed us to examine the effect of the eddies on the mean flow by calculating the barotropic E-vectors. The
 399 E-vectors indicate that the mean flow is weakened by the action of eddies, which may go some way towards explaining why vortices
 400 are shed at the highest rotation rates but not at lower rotation. Rossby wave stability theory suggests that the underlying instability
 401 leading to vortex shedding may itself be baroclinic in character.

402 It is not possible to compare our results directly with the analyses produced by Ravela *et al.* (2010), because they only compared their
 403 analyses with the observations used for assimilation (in-sample error), rather than independent observations (out-of-sample error)[†]. A
 404 comparison of the in-sample error is possible, however. Their analyses have an in-sample residual velocity error about 1.25 times
 405 the observational error (their Fig. 10). In this work we obtained similar numbers; for u the in-sample error is one to two times the
 406 observational error, and it is 1.5 to three times observational error for v . They estimated the maximum speeds in their observations to
 407 be about 2 cm s^{-1} , and their analyses had a residual velocity error around 0.15 cm s^{-1} (their Fig. 10), which is about 6% of the largest
 408 observed velocities. By way of comparison (although note that the two assimilation methods are different so a direct comparison is not
 409 possible), our assimilations have residual velocity magnitude error around $0.015\text{--}0.040 \text{ cm s}^{-1}$, which is 4–10% of a typical observed
 410 maximum velocity of 0.4 cm s^{-1} . Our method is 4–5 times slower but only uses one processor instead of four. The limiting step is
 411 to calculate the time-dependent data density, which would scale well using a multi-processor algorithm, so this could certainly be
 412 improved.

413 We aimed to demonstrate analysis correction as an example of a relatively simple meteorological technique that could be applied
 414 under laboratory conditions. In operational practice analysis correction has been superseded by newer methods, but this demonstration
 415 lays the foundations for a framework using the annulus as a tool for testing newer methods under development.

416 Newer methods that might benefit from study using the annulus include the variational methods 3D-Var (Lorenc *et al.* 2000), 4D-Var
 417 (Rawlins *et al.* 2007), and statistical 4D-Var (Lorenc and Payne 2007). The 4D- methods would require a tangent linear and adjoint
 418 model for MORALS, which we do not currently have, but we understand that one may now be available[‡]. The ensemble Kalman
 419 filter (EnKF) (Evensen 1994), particle filters (van Leeuwen 2010), and gradient descent (Stemler and Judd 2009) are other methods of
 420 interest. We have begun some work on gradient descent, which will be reported elsewhere. Zhang and Snyder (2007) list a number of
 421 challenges which (at that time) implementation of the EnKF had not yet overcome, and some of these could be usefully addressed in
 422 the context of the annulus. Model error and bias are major challenges for the EnKF because the model is used to propagate the error
 423 covariances forward in time. It might be fruitful to study this in the annulus, where specific sources of model error might be more easily
 424 identified than in atmospheric models. Ravela *et al.* (2010) use the EnKF in their annulus assimilation work, so their method might be
 425 a good place to start.

426 With these newer methods several stricter tests of the method would also be interesting. For example, we could turn off the
 427 background transform step at the start of the assimilation and see how long it takes the assimilation to push the model and observed
 428 waves into phase. Similarly, one could use an initial background state with a different dominant wave to the observations; how long does
 429 the assimilation take to converge when a wavenumber transition is required? We could also test the assimilation's ability to produce
 430 dynamically coherent structures in regions with no observations, by assimilating observations from only one part of the tank, and then
 431 verifying against the observations that were omitted. For verification, concurrent temperature measurements would be ideal, as these
 432 test both the assimilation itself and the ability of the model to recover non-observed variables, but these cannot be taken using the
 433 current experimental apparatus.

434 The setup bears some resemblance to Observing System Simulation Experiments (Arnold and Dey 1986, for example), which use
 435 a simulation combined with a model of observational error to generate artificial observations. These are then used to evaluate the
 436 performance of new proposed observation networks for NWP, for example, or to evaluate data assimilation systems. Here we have only
 437 used real data but it would certainly be possible to use simulated MORALS data in this way for experiments comparing assimilation
 438 systems, or to investigate the effects of including or omitting certain types of observation. Indeed, we have already used this 'perfect
 439 model' paradigm with MORALS in work involving breeding vectors (Young and Read 2008a).

440 Methods like the EnKF would also allow us eventually to cast the comparison of analyses with observations in a more robust
 441 statistical framework. Because the background error covariance is not updated with the analysis cycle, the current analyses represent
 442 means of distributions but the comparison is with a single realisation of the possible true states of the system. It would be better to
 443 compare the mean of the analysis distribution with the mean of the distribution of possible observed states, or ideally the comparison
 444 would be between the distribution of possible analysis states and the distribution of possible observed states. The distribution of analysis
 445 states could be obtained using the EnKF, but obtaining the observed distribution would require considerable laboratory time.

446 Acknowledgements

447 We thank Chris Farmer for some insightful comments on the results, and David Andrews, Stephen Lewis, and three anonymous
 448 reviewers for their comments. RMBY acknowledges financial support from NERC Studentship NER/S/A/2005/13667. Some figures
 449 were produced using Kevin Kohler's PSPLOT Fortran library and David Fanning's IDL library.

450 References

- 451 Arnold CP, Dey CH. 1986. Observing-Systems Simulation Experiments: Past, Present, and Future. *B. Am. Meteorol. Soc.* **67**: 687–695.
 452 Barnes SL. 1964. A Technique for Maximizing Details in Numerical Weather Map Analysis. *J. Appl. Meteorol.* **3**(4): 395–409.
 453 Bell MJ. 1989. Theoretical investigations prompted by experiments with baroclinic fluids. PhD thesis, Imperial College London.
 454 Berghórsson P, Döös BR. 1955. Numerical Weather Map Analysis. *Tellus* **7**(3): 329–340.
 455 Daley R. 1991. *Atmospheric Data Analysis*. Cambridge Atmospheric and Space Science Series, Cambridge University Press, 1st edn, ISBN 0521382157. 471
 456 pages.
 457 Dalziel S. 1995. DigImage: Image Processing for Fluid Dynamics - Particle Tracking. URL <http://www.damtp.cam.ac.uk/lab/digimage/>. 28
 458 pages.
 459 Deininger RC. 1982. Free Rossby wave instability at finite amplitude. *J. Atmos. Sci.* **39**(3): 563–572.
 460 Evensen G. 1994. Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics. *J. Geophys.*
 461 *Res. - Oceans* **99**(C5): 10 143–10 162.

[†]Sai Ravela, personal communication.

[‡]Ulrich Achatz, personal communication.

- 462 Farnell L, Plumb RA. 1976. "Numerical integration of flow in a rotating annulus II: three dimensional model". Occasional Note Met O 21 76/1, UK Met Office.
463 Unpublished, 24 pages.
- 464 Fultz D. 1951. *Compendium of Meteorology*, ch. Experimental analogies to atmospheric motions. American Meteorological Society, Boston, MA, pp. 1235–
465 1248.
- 466 Galmiche M, Sommeria J, Thivolle-Cazat E, Verron J. 2003. Using data assimilation in numerical simulations of experimental geophysical flows. *Comptes*
467 *Rendus Mecanique* **331**(12): 843–848.
- 468 Gilchrist B, Cressman GP. 1954. An Experiment in Objective Analysis. *Tellus* **6**(4): 309–318.
- 469 Gill AE. 1974. The stability of planetary waves on an infinite beta plane. *Geophys. Fluid Dyn.* **6**: 29–47.
- 470 Grotjahn R. 1984. Baroclinic instability in a long wave environment. Part I: Review. *Q. J. Roy. Meteorol. Soc.* **110**(465): 663–668.
- 471 Hide R, Mason PJ. 1975. Sloping convection in a rotating fluid. *Adv. Phys.* **24**(1): 47–100.
- 472 Hignett P, White AA, Carter RD, Jackson WDN, Small RM. 1985. A comparison of laboratory measurements and numerical simulations of baroclinic wave
473 flows in a rotating cylindrical annulus. *Q. J. Roy. Meteorol. Soc.* **111**(467): 131–154.
- 474 Hoskins BJ. 1973. Stability of the Rossby-Haurwitz wave. *Q. J. Roy. Meteorol. Soc.* **99**: 723–745.
- 475 Jackson WDN, Hignett P. 1984. A system for the real-time measurement of two-dimensional velocity fields. Met O 21 IR/84/1, UK Met Office. 46 pages.
- 476 James IN. 1994. *Introduction to circulating atmospheres*. Cambridge Atmospheric and Space Science Series, Cambridge University Press, 1st edn, ISBN
477 052141895X. 422 pages.
- 478 Judd K, Reynolds CA, Rosmond TE, Smith LA. 2008. The Geometry of Model Error. *J. Atmos. Sci.* **65**: 1749–1772.
- 479 Kalnay E. 2003. *Atmospheric Modeling, Data Assimilation and Predictability*. Cambridge University Press, ISBN 0521796296. 341 pages.
- 480 Kim K. 1978. Instability of baroclinic Rossby waves; energetics in a two-layer ocean. *Deep-Sea Res.* **25**(9): 795–814.
- 481 Lewis SR, Read PL, Collins M. 1996. Martian atmospheric data assimilation with a simplified general circulation model: orbiter and lander networks. *Planet.*
482 *Space Sci.* **44**: 1395–1409.
- 483 Lorenc AC, Ballard SP, Bell RS, Ingleby NB, Andrews PLF, Barker DM, Bray JR, Clayton AM, Dalby T, Li D, Payne TJ, Saunders FW. 2000. The Met. Office
484 global three-dimensional variational data assimilation scheme. *Q. J. Roy. Meteorol. Soc.* **126**(570): 2991–3012.
- 485 Lorenc AC, Bell RS, Macpherson B. 1991. The Meteorological Office analysis correction data assimilation scheme. *Q. J. Roy. Meteorol. Soc.* **117**(497): 59–89.
- 486 Lorenc AC, Payne T. 2007. 4D-Var and the butterfly effect: Statistical four-dimensional data assimilation for a wide range of scales. *Q. J. Roy. Meteorol. Soc.*
487 **133**(624): 607–614.
- 488 Lorenz EN. 1963. Deterministic nonperiodic flow. *J. Atmos. Sci.* **20**: 130–141.
- 489 Montabone L, Lewis SR, Read PL, Hinson DP. 2006. Validation of martian meteorological data assimilation for MGS/TES using radio occultation measurements.
490 *Icarus* **185**(1): 113–132.
- 491 Panofsky RA. 1949. Objective weather-map analysis. *J. Meteorol.* **6**(6): 386–392.
- 492 Ravela S, Marshall J, Hill C, Wong A, Stransky S. 2010. A realtime observatory for laboratory simulation of planetary flows. *Exp. Fluids* **48**(5): 915–925.
- 493 Rawlins F, Ballard SP, Bovis KJ, Clayton AM, Li D, Inverarity GW, Lorenc AC, Payne TJ. 2007. The Met Office global four-dimensional variational data
494 assimilation scheme. *Q. J. Roy. Meteorol. Soc.* **133**: 347–362.
- 495 Read PL. 1993. Coherent baroclinic waves in a rotating, stably-stratified fluid and transitions to disordered flow. In: *Waves and Turbulence in Stably Stratified*
496 *Flow*, Mobbs SD, King JC (eds). Clarendon Press, Oxford, pp. 243–274.
- 497 Read PL, Bell MJ, Johnson DW, Small RM. 1992. Quasi-periodic and chaotic flow regimes in a thermally driven, rotating fluid annulus. *J. Fluid Mech.* **238**:
498 599–632.
- 499 Sasaki Y. 1970. Some basic formalisms in numerical variational analysis. *Mon. Weather Rev.* **98**(12): 875–883.
- 500 Stemler T, Judd K. 2009. A guide to using shadowing filters for forecasting and state estimation. *Physica D* **238**(14): 1260–1273.
- 501 Thivolle-Cazat E, Sommeria J, Galmiche M. 2005. Baroclinic instability of two-layer vortices in laboratory experiments. *J. Fluid Mech.* **544**: 69–97.
- 502 Thompson PD. 1961. A Dynamical Method of Analyzing Meteorological Data. *Tellus* **13**(3): 334–349.
- 503 van Leeuwen PJ. 2010. Nonlinear data assimilation in geosciences: an extremely efficient particle filter. *Q. J. Roy. Meteorol. Soc.* **136**(653): 1991–1999.
- 504 Young RMB, Read PL. 2008a. Breeding and predictability in the baroclinic rotating annulus using a perfect model. *Nonlinear Proc. Geophys.* **15**: 469–487.
- 505 Young RMB, Read PL. 2008b. Flow transitions resembling bifurcations of the logistic map in simulations of the baroclinic rotating annulus. *Physica D* **237**(18):
506 2251–2262.
- 507 Zhang F, Snyder C. 2007. Ensemble-based data assimilation. *B. Am. Meteorol. Soc.* **88**(4): 565–568.